



HAL
open science

Intergenerational social mobility : measurement, mechanisms and policy

Gustave Kenedi

► **To cite this version:**

Gustave Kenedi. Intergenerational social mobility : measurement, mechanisms and policy. Economics and Finance. Institut d'études politiques de paris - Sciences Po, 2023. English. NNT : 2023IEPP0033 . tel-04459286

HAL Id: tel-04459286

<https://theses.hal.science/tel-04459286>

Submitted on 15 Feb 2024

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Institut d'études politiques de Paris
ÉCOLE DOCTORALE DE SCIENCES PO
Programme doctoral en économie
Département d'économie
Doctorat en sciences économiques

Intergenerational Social Mobility

Measurement, Mechanisms, and Policy

Gustave KENEDI

Thesis supervised by

Pierre-Philippe COMBES, Directeur de recherche, CNRS, Sciences Po

to be defended on November 21st, 2023

Jury

Pierre CAHUC, Professeur des Universités, Sciences Po

Gabrielle FACK, Professeure des Universités, Université Paris Dauphine - PSL (*reviewer*)

Lindsey MACMILLAN, Professor of Economics, UCL (*reviewer*)

Roland RATHELOT, Associate Professor of Economics, ENSAE-CREST

Acknowledgements

I am immensely grateful for the support and companionship I have benefited throughout my four-year journey.

I owe an immeasurable debt of gratitude to my advisor, Pierre-Philippe Combes, without whom this dissertation would never have come to fruition. Our journey began in 2016 when I was a master's student searching for a potential masters thesis advisor. Three years later, I had no doubt that I wanted to pursue a PhD under his guidance. I could not have hoped for a more caring, attentive, passionate and ever so optimistic advisor than Pierre-Philippe. Right from the beginning, he regarded me as a peer, encouraging the pursuit of my research ideas while always remaining available for my questions. During moments of self-doubt, he provided the reassurance I needed and reminded me to always celebrate even the smallest victories. Thank you, Pierre-Philippe, for your unwavering support and for being such a genuinely kind person. I hope one day we will get to work together, though I may need to enhance my theory skills for that to happen so I guess maybe not ;)

I am also profoundly thankful to the members of my dissertation committee who accepted to read my work and provide valuable comments. Pierre, your encouragements over the past years and your availability have been particularly precious, and your encyclopedic knowledge of economics has been a great source of inspiration. I extend my sincere appreciation to my dissertation's reviewers, Gabrielle Fack and Lindsey Macmillan, for agreeing to review my work at both its final stage and earlier stages. I'm also thankful to Roland Rathelot for being part of my jury.

I must express a heartfelt thank you to my two outstanding co-authors, Louis and Nagui. Louis, I vividly recall our first meeting at a café near Saint-Placide, where I hoped you'd agree that joining forces was the best approach given the proximity of our research. Our collaboration over the past three years has been a delight, and it has transformed into a cherished friendship. Given our shared meticulousness, especially regarding the details

that matter least, it's almost a surprise that we managed to complete our paper. I genuinely hope this is the first of many collaborations. Pascale, our shadow co-author and partner in crime during coffee and lunch breaks, deserves our thanks for tolerating our nerdiness and for selflessly listening to us ramble about our favorite acronyms *TSTSLS*, *IGE*, *RRC* and *PSID*. Nagui, I'm thrilled that we finally had the opportunity to work together after all those years of saying we should. Our time at the OECD was a pleasure, and our collaboration over the past year has been even more rewarding. Thank you for remaining level-headed when I had doubts, for your unwavering positivity, and for your enduring friendship.

A crucial element in any PhD student's life is the environment in which they work, and Sciences Po provided the ideal setting. I would be remiss if I didn't acknowledge the genuinely extraordinary community in "*Les Combles*." Few places concentrate so much brilliance, friendliness and humility. Without them, these four years would have lacked both meaning and, more importantly, viennoiseries. Special thanks to Naomi for your incredible generosity, kindness, friendship and for making everyone's PhD experience so much more enjoyable; Ségal for being there when I needed it the most and for your thoughtful encouragements; Mylène for being a fellow member of the PPhC fan club and for following me everywhere I go (see you in London next year!); Sophie for your warm welcome when I first joined and apologies for all the coffee breaks in my first year; Zydne for your delicious and delightful dinners (do text me when you open your gourmet kitchen!). I was fortunate to be part of a fiendly cohort of accomplished economists including Clémence, Daniel, Marco, Samuel and Victor S-J. Big thanks to my wonderful PhD tennis partners, Pierre, Juan, and Niklas, for helping me stay fit and awake at 8 a.m.

My fellow journeymen and journeywomen have made these four years immensely enjoyable: Agathe, Clemens, Claudius, Constance, Daniel, Diego, Edgard, Fatimetou, Felipe, Johanna, Leonard, Nicolas, Mattis, Moritz, Nathaniel, Nourhan, Olivia, Sophia, Rémi, Riddhi, Valentin, Venance, and Victor A. I'm also grateful to past PhD students whose journeys crossed mine, and whose imparted wisdom has influenced me: Alais, Aseem, Camille, Elisa, Jérôme, Julia, and Ludovic. My time at Sciences Po also allowed me to cross paths with Stéphane, Laurène, and Federica. Valuable advice and insights from numerous faculty members at Sciences Po have guided me, especially Ghazala Azmat, Mosche Buchinsky, Clément de Chaisemartin, Roberto Galbiati, Isabelle Méjean, Florian Oswald, and Stefan Pollinger.

I express my gratitude to Sandrine Le Goff and Pilar Calvo for their availability and diligent assistance with administrative matters. Guillaume Sarrat de Tramezaigues deserves

special thanks for ensuring the smooth day-to-day aspects of our PhD program and for his financial flexibility (and for his delightful Easter bunny impersonations).

In my fourth year, I was particularly fortunate to receive funding from PSE's *Education Policies and Social Mobility Chair*, and for that, I sincerely thank Julien Grenet and Luc Behaghel. Their financial support, along with their guidance and feedback, were instrumental during my final year. This period also allowed me to grow closer to my fellow education economists at PSE, including Amory, Cécile, Desi, Vivien, and Youssef, whose coffee breaks at PSE or DEPP were always appreciated. Antoine Bozio and David Margolis deserve my thanks for allowing me to present in their especially valuable seminar. As I write these lines, I find myself in a new job, and I extend my sincere gratitude to Sandra McNally at the Centre for Economic Performance for offering me this opportunity and for graciously allowing me to complete this dissertation unburdened. A special thanks also goes to LIEPP, and especially its director, Anne Revillard, and staff members Samira Jebli and Latifa Louaso. They provided much appreciated financial support early on.

Others have played significant roles in my PhD journey. A special thanks to Clément Dherbécourt for introducing me to the EDP, for your very generous help during my masters thesis and for transmitting your passion about intergenerational mobility; Stéphanie Jamet and Andreea Minea-Pic for encouraging me to do a PhD and for being role model researchers; Per Engzell, thank you for believing in my first paper, and for the conversations throughout my PhD; Alain Trannoy, your presence on the comité de suivi was greatly appreciated; Niels Planel for your gracious enthusiasm for my work; and José Montalbán Castilla for your detailed feedback on my third paper.

I'm also grateful to the joyful team at UCL's *Centre for Education Policy and Equalising Opportunities*, who very kindly hosted me for a few months in the spring of 2022 and introduced me to their Thursday drinks tradition. Special thanks to Jake Anders, Asma Benhenda, Stuart Campbell, Claire Crawford, Laura Outhwaite, and Gill Wyness.

This dissertation would not have been possible without access to the rich administrative data stored at the French Ministry of Education and Higher Education. I extend my gratitude to Fabienne Rosenwald, Claire Margaria, Soazig Jolivet, Axelle Charpentier, Clotilde Lixi, Cosima Bluntz, Hery Papagiorgiou, and Diane Marlat.

Of course, I am also incredibly fortunate to have my second family - my (non-PhD) friends. I don't know what I would do without you, but I am certain that my life would be less exciting and joyful. Thanks to Tiennot for all the laughs and all the tears; Simon for your sense of humour and for not thanking me in your own dissertation; Olivia for

trying to convince Simon and I that Grey's Anatomy is better than research; Alice for our shared econ nerdiness; Aldo for sharing a bed (and all your secrets) in East Ham; Ashanti for the animated political debates back in 2012; Axel and Géraldine for the delicious vegan dinners; Axelle for chats about the meaning of life until the early mornings; Céliane for imparting your love of plants; Clément and Laure (and Gustave!) for being the best of neighbors and for giving your son the best of first names; Dali for introducing me to Sergio and Samantha Power; Kamilah for all the memories in London, Copenhagen, Toronto, Berlin and Paris; Léa for your precious friendship; Mélina for the *real* conversations; Nicolas for being the best concert partner since 2010; Pauline for making sure I got the best spot in the Combles; Pauline and Pierre for your generosity and for the good times in London over the years; Sylwia for the delightful monthly conversations (as always); and Tobias for (sometimes) tolerating my terrible singing of Hamilton. I also want to thank Adrian, Amaury, Bernhard, Caro, Claire, David, Guillaume, Louis C., Manon, Madeleine, Nadège, Thibaut, and Thomas; and my flatmates, Nicolas, Guillaume, and Valentina. I love you all very dearly!

An immeasurable thank you to my family for their unconditional love and support throughout my journey, and my life more generally. To my dear father, for being so proud of me being a "Professor" at Sciences Po, even though I was really just teaching one class. To my dear mom, sorry for not ending up in the United States despite your frequent reminders. To my stepfather, who transmitted his love of problem-solving. To my brother, I don't call often enough but I treasure our lifelong bond. To my favorite aunt Ann, for being such a role model.

Last, but not least, merci Margaux pour ton soutien sans faille (et ton humour désopilant!) sur la fin de ce parcours.

London
24 September, 2023

Note to the Reader

The four chapters of this dissertation are self-contained research articles and can be read separately. They are preceded by an introduction which summarizes the research presented in this dissertation. The terms “paper” or “article” are used to refer to chapters. Chapter 1 and 2 are co-authored, which explains the use of the “we” pronoun.

Contents

Introduction	21
1 Intergenerational Income Mobility in France: A Comparative and Geographic Analysis	43
1 Introduction	43
2 Measuring Intergenerational Mobility	48
2.1 Main Measures	48
2.2 Main Sources of Bias	50
3 Data	50
3.1 Parent Income Prediction	51
3.2 Sample Definitions	55
3.3 Variable Definitions	57
4 Results at the National Level	58
4.1 Intergenerational Income Elasticity (IGE)	59
4.2 Rank-Rank Correlation (RRC)	59
4.3 Transition Matrices	63
4.4 Discussion of Baseline Results	64
5 Robustness of Baseline Results	67
5.1 Two-Sample Two-Stage Least Squares	67
5.2 Lifecycle and Attenuation Bias	68
6 Geographic Analysis	71
6.1 Heterogeneity Across Departments	71
6.2 Geographic Mobility	74
7 Conclusion	81
Appendices	87
A Data - Details	87

B	PSID Validation Exercise	89
C	Additional Robustness	115
D	Correlation with Local Characteristics	134
E	Additional Figures	140
F	Additional Tables	155
2	Older Schoolmate Spillovers on Higher Education Choices	167
1	Introduction	167
2	Institutional Background and Data	173
2.1	Institutional Background	174
2.2	Data	175
3	Empirical Strategy	175
3.1	Running Variable	176
3.2	Estimation Sample	177
3.3	Empirical Specification	180
3.4	Identifying Assumptions	182
4	Main Results	183
4.1	Within-High School Spillovers on Applications	184
4.2	Within-High School Spillovers on Enrollment	186
4.3	Cross Cohort Spillovers on College and Major	189
4.4	Robustness	189
4.5	Snowball Effect	190
4.6	Heterogeneity	191
5	Mechanisms	198
5.1	Role of Teachers	198
5.2	Information vs Role Model	199
6	Conclusion	201
	Appendices	206
A	Institutional Background: Additional Details	206
B	Running Variable: Additional Details	209
C	Robustness Checks	211
D	Appendix Figures	214
E	Appendix Tables	221
3	High-Achieving, Low-Income Students and Higher Education: What Can Financial Aid Achieve?	223

1	Introduction	223
2	Institutional Background	229
	2.1 Aide au Mérite	229
	2.2 Need-Based Grants	230
	2.3 Timeline	231
	2.4 Higher Education in France	232
3	Data, Sample Construction and Descriptive Statistics	234
	3.1 Data	234
	3.2 Sample	235
	3.3 Descriptive Statistics	235
4	Empirical Strategy	236
	4.1 Estimation Details	236
	4.2 Tests of Design Validity	240
5	Main Results	242
	5.1 Academic Outcomes: Enrollment, Degree Quality, Persistence, Grad- uation, and Academic Performance	243
	5.2 Non-Academic Outcomes: Geographic Mobility	249
	5.3 Heterogeneity	251
6	Mechanisms	253
	6.1 Were Students Unaware of the Policy?	253
	6.2 Did the Aid Crowd Out Parents' Financial Support?	253
	6.3 Was the Awarded Amount Too Small Relative to Need-Based Grants?	255
	6.4 Discussion: Were Targeted Students Not Marginal Students?	257
7	Conclusion	258
	Appendices	264
	A Appendix Figures	264
	B Appendix Tables	275
	C Detailed Regression Tables	280
	Résumé en Français	303

List of Figures

1.1	OLS vs. TSTSLS RRC - National and Census Regions in the United States . . .	54
1.2	Out-of-Sample Predicted Labor Income Rank - France (EDP) and United States (PSID)	55
1.3	Conditional Expectation Functions for Log-Log and Rank-Rank Relationships in France	60
1.4	Baseline Quintile Transition Matrix for Different Countries	63
1.5	Graduation From/Enrollment In Higher Education by Parent Income	66
1.6	Child Lifecycle Bias	69
1.7	Parent Lifecycle Bias	70
1.8	Spatial Variations in Intergenerational Mobility	73
1.9	Intergenerational Mobility and Department Characteristics - Separate Estimation	75
1.10	Intergenerational Mobility and Geographic Mobility	76
1.11	Mean Income Rank of Origin and Destination Departments of Movers	78
1.12	Mean Child Income Rank by Destination Department Mean Income	80
B.1	Benchmark Rank-Rank Correlation and Transition Matrix	92
B.2	Benchmark Rank-Rank Correlation and Absolute Upward Mobility by Census Region	93
B.3	OLS vs. TSTSLS RRC and Transition Matrix	95
B.4	OLS vs. TSTSLS RRC and AUM	96
B.5	Benchmark IGEs for All Income Definitions	100
B.6	Benchmark RRCs for All Income Definitions	100
B.7	Benchmark Rank-Rank Correlation and Transition Matrix	101
B.8	Observed vs. (out-of-sample) predicted <i>individual</i> labor income	102
B.9	Observed vs. (out-of-sample) predicted <i>individual</i> labor income rank	103
B.10	Observed vs. (out-of-sample) predicted <i>individual</i> labor income quintile . .	104
B.11	Observed vs. (out-of-sample) predicted <i>family</i> labor income	105

B.12	Observed vs. (out-of-sample) predicted <i>family</i> labor income rank	106
B.13	Observed vs. (out-of-sample) predicted <i>family</i> labor income rank	107
B.14	OLS vs. TSTSLS Estimates - Varying Number of Parent Income Observations	111
B.15	OLS vs. TSTSLS Estimates - Varying Number of Parent Income Observations - Child Income Mean 37-43	112
C.1	Robustness of Baseline Estimates to Computing Synthetic Parent Incomes only on Post-1988 Data	115
C.2	Robustness of Baseline Estimates to Different First-Stage Predictors	119
C.3	Robustness to Machine Learning Prediction	121
C.4	Child Lifecycle Bias - 1972 and 1974 Cohorts (Constant Sample)	122
C.5	Child and Parent Lifecycle Bias	123
C.6	Attenuation Bias	124
C.7	Parent Attenuation Bias (Constant Sample of Synthetic Parents)	126
C.8	Sensitivity to Number of Child Income Observations (Constant Sample)	127
C.9	Sensitivity to Different Zero Child Income Replacement Values	129
C.10	Sensitivity to Child and Parent Income Distributions Trimming	131
C.11	Out-of-sample MSE as a Function of Top and Bottom Trimming	132
C.12	Top Parent Income Quantiles Transition Matrices in France and United States	133
D.1	Correlation Between Department Characteristics	135
D.2	Department Characteristics Kept by Lasso	136
E.1	Family Position in 1990 Census by Child Birth Cohort	140
E.2	Baseline IGE Estimates for Different Child and Parent Income Definitions	141
E.3	Baseline RRC Estimates for Different Child and Parent Income Definitions	142
E.4	Baseline Quintile Transition Matrix for Different Child Income Definitions	143
E.5	Rank-Rank Correlation and Upward Mobility in International Comparison	144
E.6	Higher Education Graduation by Quintile Transition Matrix Cell	145
E.7	French Departments	146
E.8	Illustration of Absolute Upward Mobility for the <i>Nord</i> Department	147
E.9	Department-Level Log-Log Relationships	148
E.10	Department-Level Rank-Rank Relationships	149
E.11	Department-Level Intergenerational Elasticities	150
E.12	Department-Level Rank-Rank Correlations	151
E.13	Department-Level Absolute Upward Mobility	152
E.14	Department-Level Unemployment Rate in 1990	153
E.15	Geographic Mobility by Parent Household Wage Rank	153
E.16	Intergenerational Mobility and Geographic Mobility - Department Ranks	154

2.1	Probability of Older Schoolmate Enrolling in College-Major as a Function of the Distance to the Last Admitted Student	181
2.2	Distribution of Running Variable	183
2.3	Discontinuity in High School Characteristics	184
2.1	Within-High School Spillovers on Applications to and Enrollment in College-Major, College and Major with Marginally Admitted Older Schoolmate . . .	187
2.2	Placebo and Snowball Effects	191
2.3	Heterogeneity by Student Characteristics	193
2.4	Heterogeneity by High School Characteristics	194
2.5	Heterogeneity by College-Major Characteristics	196
2.6	Heterogeneity by High School <i>and</i> College-Major Characteristics	197
2.1	Older Schoolmate Teacher-Specific Spillovers	199
2.2	Older Schoolmate Gender- and SES-Specific Spillovers	200
A.1	Higher Education Landscape in France	207
A.2	Timeline of the Application and Admission Procedure into Higher Education Programs in France	208
B.1	An Example of the Running Variable for One College-Major	210
C.1	Robustness of Baseline Within-High School Spillovers to Using Different Bandwidths	212
C.2	Robustness of Baseline Within-High School Spillovers to Placebo Admission Cutoffs	213
D.1	Exact Number of Enrolled Students in Treatment Year	214
D.2	Running Variable With Degree Type Composition	214
D.3	Discontinuity in High School Characteristics	215
D.4	Discontinuity in College-Major Characteristics	216
D.5	Event-Study Analysis Graphs	217
D.6	Effects for College-Majors with No Applicants from High School in $t - 1$. .	218
D.7	Role Model Effects for Girls	218
D.8	Role Model Effects for Boys	219
D.9	Role Model Effects for Low SES Students	219
D.10	Role Model Effects for Very High SES Students	220
3.1	Timeline of Events	232
3.1	Distribution of Bac Grades, 2009-2014	239
3.1	Higher Education Outcomes as a Function of Bac Grade Around the Aide au Méri-te Eligibility Threshold (16/20)	246

3.2	Higher Education Outcomes as a Function of Bac Grade Around the Aide au Mérite Eligibility Threshold (16/20)	250
3.3	Heterogeneity of Effects of Eligibility to the Aide au Mérite in Bac Year by Gender, SES, and High School Track	252
A.1	Aide au mérite Status Over Time	264
A.2	Need-Based Grant Eligibility Over Time	265
A.3	Distribution of Bac grades, 2009-2014, all students	265
A.4	Distribution of Bac grades, 2009-2014 - 16.05 Grades In [16, 16.05) Bin	266
A.5	Discontinuity Estimates on Observable Characteristics for Different Donut Limits	267
A.6	Discontinuity in Pre-Treatment Observable Characteristics	268
A.6	Discontinuity in Pre-Treatment Observable Characteristics (<i>continued</i>)	269
A.7	Effect of Eligibility to the Aide au Mérite in Bac Year on Degree Choice	270
A.8	Effect of Eligibility to the Aide au Mérite on Obtaining a Degree	271
A.9	Effect of Eligibility to the Aide au Mérite on Location of Studies	272
A.10	Effect of Eligibility to the Aide au Mérite in Bac Year on Higher Education Outcomes by Bandwidth Size	273
A.11	Effect of Eligibility to the Aide au Mérite on Location of Studies by Bandwidth Size	274

List of Tables

1.1	Rank-Rank Correlation in International Comparison	62
1.2	Transition Matrix in International Comparison	65
1.3	Intergenerational & Geographic Mobility	77
B.1	Descriptive Statistics	91
B.2	Sample Size at Each Restriction	97
B.3	Detail of Variables Used in Each Year	99
B.4	Comparison for Different Sets of Predictors	108
B.5	Robustness of Baseline Results	110
B.6	Comparison between Baseline Results and Weighted Results	114
C.1	Average Characteristics of Overall Population vs. Sample	117
D.1	Definitions and Sources of Department Characteristics	134
D.2	Correlation Between Intergenerational Elasticity and Department Characteristics	137
D.3	Correlation Between Rank-Rank Correlation and Department Characteristics	138
D.4	Correlation Between Absolute Upward Mobility and Department Characteristics	139
F.1	Child Sample Construction	155
F.2	Average Characteristics of Actual and Synthetic Parents	156
F.3	Share of Actual and Synthetic Parents by 2-Digit Occupation	157
F.4	Synthetic Parents Sample Construction	158
F.5	Number of Observations by Child and Parent Income Definitions	158
F.6	Descriptive Statistics	159
F.7	Comparison with Existing Father-Son IGE Estimates for France	160
F.8	Department-Level MSEs and Measures of Intergenerational Income Mobility	160
F.9	Department-Level Intergenerational Mobility Estimates	161
F.10	Correlation Between Department-Level Intergenerational Mobility Measures	164
F.11	Intergenerational & Geographic Mobility - Department Ranks	164

F.12	Intergenerational Mobility and Income Level in the Destination Department	165
2.1	Descriptive Statistics	179
2.1	High School Cohort Spillovers on Applications to and Enrollment in Degree, College and Major with Marginally Admitted Student	188
2.2	College-Major Spillovers on Applications to and Enrollment for College-Majors with No Applications from High School in t-1	192
E.1	Number of Observations at Each Sample Restriction	221
3.1	Discontinuity Estimates for Pre-Treatment Observable Characteristics	241
3.2	Placebo Analysis	243
3.1	Effects of Eligibility to the aide au merite in Bac Year on Higher Education Outcomes	244
3.2	Effects of Eligibility to the Aide au Mérite in Bac Year on Degree Choice	247
3.3	Effects of Eligibility to the Aide au Mérite in Bac Year on Location of Studies	251
3.1	Effects of Eligibility to the Aide au Mérite on Enrollment in Bac Year by Student Awareness Proxy	254
3.2	Test of Crowding Out Effects of Aide au Mérite Eligibility	256
3.3	Test of Whether aide au mérite Eligibility Had Effects for Students with No or Small Amounts of Cash Allowance as Part of their Need-Based Grant	257
3.4	Summary of Studies Finding Complementarities Between Financial Aid and Academic Ability	258
B.1	Combinations of Parent Income and Disadvantage Points for each Need-Based Grant Echelon for the 2009-10 Academic Year	275
B.2	Need-Based Grants Annual Amounts by Echelon	275
B.3	Number of Observations at Each Sample Restriction	276
B.4	Descriptive Statistics	277
B.5	Discontinuity in Predicted Outcomes at Aide au Mérite Eligibility Threshold (16/20 Bac Grade)	278
B.6	Placebo Analysis using (15, 17) Bandwidth	279
B.7	Test of Whether aide au mérite Eligibility Had Effects for Students with No or Small Amounts of Cash Allowance as Part of their Need-Based Grant - Parent Income Within 5% of Parent Income Threshold (2009-2012)	279
C.1	Effect of Eligibility to the Aide au Mérite in Bac Year on Enrollment in Bac Year	280
C.2	Effect of Eligibility to the Aide au Mérite in Bac Year on Enrollment at Least Once	281

C.3	Effect of Eligibility to the Aide au Mérite in Bac Year on Enrollment in 2 nd Year in Bac Year + 1	282
C.4	Effect of Eligibility to the Aide au Mérite in Bac Year on Enrollment in 2 nd Year at Least Once	283
C.5	Effect of Eligibility to the Aide au Mérite in Bac Year on Enrollment in 3 rd Year in Bac Year + 2	284
C.6	Effect of Eligibility to the Aide au Mérite in Bac Year on Enrollment in 3 rd Year at Least Once	285
C.7	Effect of Eligibility to the Aide au Mérite in Bac Year on the Number of Years Enrolled in Higher Education	286
C.8	Effect of Eligibility to the Aide au Mérite in Bac Year on the Highest Level of Study Attained (in Years)	287
C.9	Effect of Eligibility to aide au mérite on Degree Quality in Bac Year	288
C.10	Effect of Eligibility to aide au mérite on Degree Quality in Bac Year + 1	289
C.11	Effect of Eligibility to aide au mérite on Degree Quality in Bac Year + 2	290
C.12	Effect of Eligibility to the Aide au Mérite in Bac Year on Obtaining a Degree	291
C.13	Effect of Eligibility to the Aide au Mérite in Bac Year on Enrolling in a Masters Degree	292
C.14	Effect of Eligibility to the Aide au Mérite in Bac Year on Enrolling in a Selective Masters Degree	293
C.15	Effect of Eligibility to aide au mérite on First Graduate Degree Quality	294
C.16	Effect of Eligibility to the Aide au Mérite in Bac Year on Enrolling in Paris (Urban Unit) in Bac Year	295
C.17	Effect of Eligibility to the Aide au Mérite in Bac Year on Enrolling in Paris, Marseille or Lyon (Urban Units) in Bac Year	296
C.18	Effect of Eligibility to the Aide au Mérite in Bac Year on Enrolling in a City with Over 200k Inhabitants in Bac Year	297
C.19	Effect of Eligibility to the Aide au Mérite in Bac Year on Enrolling in Paris (Urban Unit) in Bac Year + 1	298
C.20	Effect of Eligibility to the Aide au Mérite in Bac Year on Enrolling in Paris, Marseille or Lyon (Urban Units) in Year Bac + 1	299
C.21	Effect of Eligibility to the Aide au Mérite in Bac Year on Enrolling in City with Over 200k Inhabitants in Bac Year + 1	300

Introduction

“There is no extravagance more prejudicial to the growth of national wealth than that wasteful negligence which allows genius that happens to be born of lowly parentage to expend itself in lowly work.”

– Alfred Marshall, *Principles of Economics* (1890)

“The extent to which natural capacities develop and reach fruition is affected by all kinds of social conditions and class attitudes. Even the willingness to make an effort, to try, and so to be deserving in the ordinary sense is itself dependent upon happy family and social circumstances. It is impossible in practice to secure equal chances of achievement and culture for those similarly endowed, and therefore we may want to adopt a principle which recognizes this fact and also mitigates the arbitrary effects of the natural lottery itself.”

– John Rawls, *A Theory of Justice* (1971)

TO what extent are individuals’ life outcomes shaped by their childhood circumstances, such as their parents’ incomes, the neighborhood(s) they grew up in, their school teachers, etc.? In the United States, only 7.5% of children born in the early 1980s to parents in the bottom 20% of the income distribution reach the top 20% as adults (Chetty et al., 2014). In a society with no relationship between parent and child income this probability would be 20%, since one’s income quintile of origin would be unrelated to one’s future outcomes. What explains such a strong persistence in incomes across generations? What policies may help remediate these intergenerational inequalities? How do other countries compare with the United States? In particular, does France, a country with significantly less income inequality and a relatively inexpensive higher education system, exhibit more intergenerational mobility?

These questions lie at the heart of my dissertation. Each chapter aims to (modestly) im-

prove our understanding of the possible answers to these questions. Chapter 1, joint with Louis Sirugue, measures the extent of persistence in incomes across generations in France, and compares the results to those found for other advanced economies. Chapter 2 aims to better understand one of the mechanisms underlying intergenerational immobility, specifically how students choose which college and major to pursue after high school. In particular, it explores how students are influenced by the higher education trajectories of older schoolmates. Chapter 3 evaluates whether higher education enrollment and graduation gaps between high-achieving, low-income students and their high-income peers can be reduced by providing these students with additional financial support.

The first chapter provides new estimates of intergenerational income mobility in France for children born in the 1970s. Surprisingly, in the country of Pierre Bourdieu, very little is known about intergenerational income mobility. Existing studies relied on small samples, self-reported incomes, focused exclusively on fathers and sons, and measured intergenerational mobility using only one statistic (Lefranc and Trannoy, 2005; Lefranc, 2018). In recent years, a very rich administrative dataset combining census and tax returns, the *Permanent Demographic Sample*, has been made available to researchers. We use it to estimate intergenerational mobility for a much larger sample and significantly better information on individuals' incomes. We incorporate the latest developments of the literature, in particular, grouping sons and daughters together, accounting for mothers by defining income at the *household* level, and estimating more recently proposed measures of intergenerational mobility such as the rank-rank correlation. Though parents' incomes continue to be unobserved, we leverage the large number of available information about parents, such as their education level, their occupation, where they live, etc. to predict their incomes using the two-sample two-stage least squares (TSTSLS) procedure.

We find that France is characterised by a strong persistence in incomes across generations relative to other advanced economies. Only 9.7% of children from the bottom 20% of the parent income distribution reach the top 20% as adults, almost 4 times less than children born to parents in the top 20% (38.4%). In comparison, the probability for a child born to a family in the bottom 20% to reach the top 20% in adulthood is 7.5% in the United States (Chetty et al., 2014) and 12.3% in Australia (Deutscher and Mazumder, 2020). We also document very large spatial variations in intergenerational mobility across French departments. They appear to be most related to the unemployment rate while growing up. Lastly, we find important gains to geographic mobility. In particular, the expected income rank of individuals from the bottom of the parent income distribution who moved towards high-income departments is around the same as the expected income rank of

individuals from the 75th percentile who stayed in their childhood department.

This paper's results only raise further questions. Why is intergenerational mobility so low in France? More generally, what determines intergenerational mobility? Why are children from high-income families more likely to have incomes themselves? Why are some countries more intergenerationally mobile than others? There are, of course, a multitude of factors that can shed light on the underlying causes of intergenerational mobility. [Bowles \(1973\)](#) categorized these determinants into three types of explanations: (i) inequality in educational opportunities, (ii) differences in aspirations, personality characteristics and other family background cultural traits, (iii) inheritability of intellectual abilities. [Chapters 2](#), which is joint work with Nagui Bechichi, and [chapter 3](#) contribute to our understanding of explanations at the intersection between inequalities in (higher) educational opportunities and differences in aspirations due to family background.

The second chapter explores the role played by older schoolmates' higher education trajectory in shaping students' higher education choices. Deciding whether to apply to college, and choosing the right college and major is a highly complex decision. Yet this decision is important: graduating from higher education provides one of the highest returns on investment an individual can make, though the returns vary across majors, and to some extent across institutions ([Altonji et al., 2012](#); [Hastings et al., 2013](#); [Kirkeboen et al., 2016](#); [Aucejo et al., 2022](#); [Black et al., 2023](#); [Chetty et al., 2023](#)). Students are unevenly equipped to make this decision, due to differences in information about the returns to higher education, knowledge of college and majors, aspirations due to family background, or simply differences in financial resources.

Recent work has highlighted the important roles played by students' social networks such as their family (parents and siblings) and close ties (neighbors, peers, teachers) ([Aguirre and Matta, 2021](#); [Altmejd et al., 2021](#); [Barrios-Fernández, 2022](#); [Altmejd, 2023](#)). This suggests exposure to peers' higher education choices could be an important source of information about higher education. In this chapter we investigate, for the first time, the extent to which, within the same high school, students' applications and enrollment choices are influenced by older schoolmates' higher education trajectories.

Using a regression discontinuity design and French administrative data on students' higher education applications, we find sizable within-high school spillovers. Students are significantly more likely to apply to and enrol in a college-major if in the previous year there was a marginally admitted student to that same college-major who was from the same high school, compared to students in high schools with a marginally rejected

student. More generally, they are more likely to apply to the same college, though we do not find any spillovers on majors. Moreover, the magnitude of these spillovers vary with high school, college-major, and the interaction between high school and college-major characteristics. Smaller high schools exhibit larger cross-cohort spillovers, as well as less selective college-majors. Geographic distance seems to play a role, with very close and semi distant college-majors inducing the largest spillovers in terms of applications. We find that student role model effects explain most of these spillovers, rather than teachers. Girls are significantly more apply to a college-major if the marginal admitted older schoolmate was a girl rather than a boy, and conversely for boys. These results highlight the important role played by students' high school environment in shaping their higher education choices.

The third chapter assesses whether increased financial assistance can mitigate the enrollment gaps observed between high-achieving, low-income students and their high-income peers. Across a number of countries, there is a large gap in enrollment, quality of institution attended and graduation between students of different socio-economic backgrounds, even after conditioning on high school achievement (Hoxby and Avery, 2013; Crawford et al., 2016; Dynarski et al., 2021; Campbell et al., 2023; Hakimov et al., 2022). What explains such large gaps? Is the explanation that high-achieving, low-income students are less aware of the benefits of attending higher education or lack information about relevant programs? Or is it that these students require additional financial resources to attend these selective colleges?

In this chapter, I estimate the impact of automatically granting additional financial support to high-achieving, low-income students who enrol in higher education. Using comprehensive administrative data for France and a regression discontinuity design, I find that this policy had no significant effect on enrollment, persistence, graduation or academic performance in higher education. I also find no evidence that this aid induced eligible students to enrol in or switch to higher quality degrees during their studies. There are two main takeaways. First, at least in the French context, additional financial support on top of existing programs, without any other changes, does not seem to have impacted any of the relevant academic margins. However, the policy could have had positive effects on students' mental health and financial distress, which are not observed in the data. Second, based on these results and those found in the literature, there appears to be complementarities between financial aid and academic ability. In particular, students with lower academic levels at the point of college entry are likely to be significantly more adversely impacted by a lack of financial support relative to students with greater college

readiness.

Below, I describe each chapter in more detail.

Chapter 1: Intergenerational Income Mobility in France

Co-authored with Louis Sirugue (Paris School of Economics)

To what extent is the income of individuals related to that of their parents? This question has seen renewed interest both in the general public and in academia as rising income inequality raised concerns about equality of opportunity. Examining this link is essential to understand whether children from different socio-economic backgrounds are afforded the same opportunities. It also matters for economic efficiency, as high persistence across generations may reflect an inefficient allocation of talents (so-called “Lost Einsteins”). Intergenerational persistence has now been estimated for a large number of countries, paving the way for insightful cross-country comparisons. Yet, much remains to be known for France, a country with relatively modest post-tax/transfers income inequality in international comparison and largely inexpensive higher education tuition fees.

The few existing studies for France only estimate the traditional intergenerational income elasticity (IGE), which captures the elasticity of child income with respect to parent income, and are based on small-sample surveys with self-reported incomes (Lefranc and Trannoy, 2005; Lefranc, 2018). Using a large sample combining census and tax returns data, we estimate two additional measures of intergenerational mobility: (i) the rank-rank correlation (RRC), increasingly prominent in the literature, which corresponds to the correlation between child and parent income percentile ranks, and (ii) transition matrices, which capture finer mobility patterns along the parent income distribution. While previous studies on France used self-reported labor earnings, we focus on household-level income measures. They provide a better depiction of one’s economic resources and allow the inclusion of children raised by single mothers. Integrating these improvements from the “new” intergenerational mobility literature enables us to conduct a detailed international comparison to rank France relative to other advanced economies for which comparable estimates are available.

In addition, we investigate the spatial variations in intergenerational mobility across the 96 metropolitan French departments. Such subnational analyses, pioneered by Chetty et al. (2014), help shed light on the mechanisms that may underlie income persistence across generations. Importantly, they highlight that national level estimates provide an

incomplete assessment of a country's intergenerational mobility. We make use of the panel dimension of our data to describe the geographic mobility patterns of individuals and study the relationship between geographic mobility and intergenerational mobility. We investigate the separate roles of moving to a higher-income department from that of climbing the income ladder within departments, conditional on parent income rank.

Our analysis is conducted on almost 65,000 children born between 1972 and 1981, and observed in the Permanent Demographic Sample (EDP). This rich administrative dataset allows us to implement the contributions discussed above and to convincingly address concerns related to lifecycle and attenuation bias (Haider and Solon, 2006; Black and Devereux, 2011; Nybom and Stuhler, 2017). Since parents' incomes are not observed, we use a two-sample two-stage least squares (TSTSLs) estimation which consists in predicting parents' incomes using other parents drawn from the same population but for whom income is observed (Björklund and Jäntti, 1997). This method has been employed previously to estimate the IGE in the French context (Lefranc and Trannoy, 2005; Lefranc, 2018) as well as in many other countries (Jerrim et al., 2016, Table A1).

While studies typically use education and/or occupation to predict parent income, we make use of the richness of our data to also include detailed demographic characteristics of parents (French nationality dummy, country of birth, household structure, and birth cohort), and characteristics of the municipality of residence (unemployment rate, share of single mothers, share of foreigners, population, and population density). Our results are largely insensitive to the set of predictors. Parent income is then defined as the average of father and mother predicted mean pretax wage over ages 35-45, and child income as pretax household income averaged over the same age range between 2010 and 2016. These two income definitions represent the most comprehensive household-level income definitions available for either generation.

TSTSLs Validation Exercise. Using the United States' Panel Study of Income Dynamics (PSID), we find that TSTSLs slightly underestimates rank-based measures of intergenerational persistence relative to what would be obtained if parent income were observed (OLS). The downward bias relative to the OLS estimate for the RRC ranges from 11% when education is the only predictor, to around 3-5% once occupation is also included. Subnational TSTSLs estimates are also fairly close to their OLS counterparts, though they tend to deviate more when the number of observations is small. Our results highlight that in settings like ours, where parent income cannot be directly observed, rank-based measures of intergenerational mobility obtained with TSTSLs likely provide lower bounds

that are reasonably close to the true estimates. These findings confirm those obtained in different settings and samples by [Cortes-Orihuela et al. \(2022\)](#) and [Jacome et al. \(2023\)](#). We find that this reasoning also applies to the transition matrix.

National Results. Our main finding is that France exhibits relatively strong intergenerational income persistence compared to other developed countries. Our baseline estimate of the intergenerational elasticity in household income is 0.527, suggesting that on average, a 10% increase in parent income is associated with a 5.27% increase in child income. Put differently, if one's parents earn 10% more than the average of parents' incomes, then one is expected to preserve about 50% of that relative advantage. This estimate should be interpreted with caution considering our validation exercise suggests the TSTOLS IGE is significantly greater than the true estimate. Applying the correction factor we find, the IGE decreases to 0.396.

Moving to the rank-rank relationship, we find that the conditional expectation of child income percentile rank with respect to parent income percentile rank is linear throughout most of the parent income distribution, with steeper relationships at the tails. Our baseline estimate of the rank-rank correlation is 0.303, implying that a 10 percentile increase in parent income rank is associated, on average, with a 3.03 percentile increase in child income rank. This estimate is of similar magnitude to that found for Italy (0.3; [Acciari et al. \(2022\)](#)), somewhat smaller than for the United States (0.341; [Chetty et al. \(2014\)](#)), and markedly greater than existing estimates for other advanced economies such as Sweden (0.197; [Heidrich \(2017\)](#)), Australia (0.215; [Deutscher and Mazumder \(2020\)](#)) or Canada (0.242; [Corak \(2020\)](#)). Applying the correction factor we find in the validation exercise gives an RRC of 0.314 which does not affect France's relative position.

Intergenerational persistence, as captured by the transition matrix, is strongest at the tails of the parent income distribution: 9.7% of children from the bottom 20% of the parent income distribution reach the top 20% as adults. This probability is almost 4 times greater for children born to parents in the top 20% (38.4%). In comparison, the probability for a child born to a family in the bottom 20% to reach the top 20% in adulthood is 7.5% in the United States ([Chetty et al., 2014](#)) and 12.3% in Australia ([Deutscher and Mazumder, 2020](#)). Moreover, persistence at the top becomes stronger and stronger as we zoom in on the right tail of the parent income distribution. As with the RRC, the validation exercise suggests these estimates represent upper (lower) bounds on mobility (persistence).

We show that our baseline results are robust to potential biases. Foremost, we evaluate

how sensitive they are to the parent income prediction specification. In particular, we check whether varying the set of predictors or using non-parametric estimation methods influences our estimates. IGE estimates are overinflated when using only education as a predictor, while the RRC and transition matrices remain surprisingly stable regardless of the set of predictors used. Slightly improved prediction from using flexible models does not quantitatively alter our estimates. Moreover, we assess our estimates' sensitivity to the lifecycle and attenuation biases by varying the ages at which child and parent incomes are measured as well as the number of parent income observations used. Our baseline results do not appear to under- nor over-estimate intergenerational mobility due to measuring child and/or parent incomes too early or too late in the lifecycle nor because of averaging incomes over too few years.

Subnational Results. We uncover substantial spatial variations in intergenerational mobility across departments, comparable to those observed across countries. We define individuals' location as their department of residence in 1990, when they are between 9 and 18 years old. Higher levels of mobility are typically found in the West of France, and lower levels in the North and South. While the IGEs range from 0.30 to 0.45 in departments in Brittany (West), they range from 0.42 to 0.70 in departments in Hauts-de-France (North). The distribution of department-level RRCs is tighter than that of IGEs, but displays very similar spatial patterns.

We also characterize departments' absolute upward mobility (AUM), defined as the expected income rank of children born to parents at the 25th percentile, which is obtained from the fitted values of the department-level rank-rank regression (Chetty et al., 2014). Absolute upward mobility ranges from the 36.8 in Pas-de-Calais (North) to 54.4 in Haute-Savoie (East). The Paris department stands out in terms of AUM (49.8) but exhibits around average intergenerational persistence levels in terms of IGE (0.51) and RRC (0.28). The cross-department correlation between the IGE and RRC is only 0.65, and -0.55 with AUM. This highlights the importance of using a variety of intergenerational mobility measures to characterize a country's income persistence across generations (Deutscher and Mazumder, forthcoming).

As a first step to understand the sources underlying these cross-department variations in intergenerational mobility, we undertake a simple correlational analysis. We find that absolute upward mobility exhibits much stronger relationships with department characteristics in general, than either the IGE or the RRC. This suggests that factors that affect absolute mobility might differ from those that affect relative mobility. The only character-

istic consistently negatively correlated with intergenerational mobility is the unemployment rate. Intriguingly, we find no evidence of a within France “Great Gatsby Curve”¹ with respect to the IGE nor the RRC. This contrasts with findings from other countries (Acciari et al., 2022; Chetty et al., 2014; Corak, 2020).

Lastly, we conduct a descriptive analysis of the relationship between intergenerational income mobility and geographic mobility. We document important gains in expected income rank for movers, which are slightly decreasing in parent income rank. For children from families in the bottom decile, movers have an expected rank approximately 5.6 percentiles greater than stayers, while this difference is of roughly 4.4 percentiles for children from families in the top decile. These gains are partly attributable to movers locating in higher-income departments in adulthood relative to stayers, but also to movers reaching local ranks in their adulthood department that are further away from the rank of their parents in the childhood department. Destination departments are on average characterized by higher income levels than origin departments only at the tails of the parent income distribution. However, regardless of parent income rank, the absolute upward mobility gains associated with moving to a higher-income department appear to be large and increasing with average income in the destination department. All these findings combine self-selection and causal effects, and we leave the disentangling of these two channels for future research.

Chapter 2: Older Schoolmate Spillovers in Higher Education Choices

Co-authored with Nagui Bechichi (Paris School of Economics)

How do students choose whether and where to apply to university? Considering the large returns to higher education, and the large differences across majors and institutions, this question has received tremendous attention. Recent work has highlighted the important roles played by informational deficits (Hoxby and Turner, 2013a; Carrell and Sacerdote, 2017), and by students’ social networks, such as their parents (Altmejd, 2023), their siblings (Aguirre and Matta, 2021; Altmejd et al., 2021) and even their neighbors (Barrios-Fernández, 2022). This suggests exposure to peers’ higher education choices might be an important source of information for students’ decisions. However, we know very lit-

¹The “Great Gatsby Curve” refers to the positive correlation between intergenerational income persistence (defined by the IGE) and income inequality (defined by the Gini index) found across countries (Corak, 2013).

tle about how high school peers shape this decision. In particular, within the same high school, how are students' applications and enrollment choices influenced by older school-mates' higher education trajectories? Causally identifying such effects is challenging due to the important sorting of students across high schools and the endogeneity of students' higher education choices to their high schools.

This paper provides the first causal evidence on within-high school² spillovers on higher education choices. Using administrative application data from France covering close to 90% of higher education programs between 2013 and 2017 (Bechichi et al., 2021), we show that students are more likely to apply to and enrol in a college-major³ if a student from the same high school enrolled in this exact same college-major the previous year. We also find important spillovers on the choice of college more broadly, but no effects on the chosen major.

We identify within-high school spillovers by exploiting admission cutoffs generated by France's centralised admission procedure. This allocation ensures programs cannot anticipate ex-ante the high school of the last admitted student. As such high schools around a college-major's admission threshold are virtually identical other than for having a student ranked just above or just below the rank of the last admitted student to this college-major. This generates quasi-random variations in the college-majors to which a high school's students are admitted to and enrol in, which in turn also generates quasi-random variations in the college-majors to which the following cohort of students in the same high school is exposed to. This enables us to implement a fuzzy regression discontinuity design to estimate within-high school spillovers on applications and enrollment. While existing work has exploited comparable cutoffs generated by *academic* thresholds in admission policies (e.g., Altmejd et al. (2021); Estrada et al. (2022)), our design is very similar in spirit except we only observe the relative *ranking* of students by the college-majors to which they have applied. Since several students from the same high school may apply to the same college-major, we keep only the high school's best ranked applicant by the college-major, as in Estrada et al. (2022).

We find that students follow the higher education choices of their high school's previous graduating cohort. They are 7 percentage points (+25% relative to the counterfactual mean) more likely to apply to, and 3 percentage points (+67%) more likely to enroll in, a college-major in which a student from their high school's previous cohort was marginally

²Technically, our analysis is undertaken at the high school x track level because, in France, high school tracks are very segregated within high schools and higher education programs are often largely track-specific. To ease legibility, we use "high school" to refer to "high school x track".

³We will use interchangeably *college-major*, *college program* and *program*.

admitted to and enrolled in, relative to students in high schools with a marginally rejected older schoolmate. We also uncover large impacts on the intensive margin, i.e., the number, of applications and enrolled students: 0.24 (+30%) and 0.05 (+72%) percentage points increases respectively.

The magnitude of these effects is large. Compared to sibling spillovers in college-major estimated by [Altmejd et al. \(2021\)](#) for Chile, Croatia and Sweden, our within-high school spillovers are between 43% and 78% as large as the impact on applications they find, and between 40% and 88% of their enrollment effects. Moreover, we also show that students are more likely to apply to and enrol in the same college as their older high school peers but there are no spillovers on major choice. The (relative) magnitude of the college spillovers are roughly similar to the college-major spillovers, 9.6 percentage points (+17%) for applications and 10.9 percentage points (+52%) for enrollments. The lack of spillovers on majors could potentially be explained by the fact that students have stronger preferences over what they want to study than over where they want to study or because they are more aware of existing majors. Therefore the college component of a previous peer's enrollment is more salient to them. This result is in line with [Altmejd et al. \(2021\)](#) and [Aguirre and Matta \(2021\)](#) who also find no sibling spillovers on majors.

We uncover several insightful heterogeneities with respect to college-major spillovers. First, we find that the magnitude of the spillovers are broadly constant over the four outcome years. This suggests they are not the result of a given year's idiosyncrasies, but rather a structural determinant of students' higher education choices. Second, with regards to student characteristics, we find that within-high school spillovers on applications are of similar magnitude for both genders, though the effects on enrollment are significantly larger for boys. This could be driven by differences in the types of degrees applied to. Moreover, and quite surprisingly, we find that low socioeconomic status (based on legal guardian's occupation) students are only slightly more responsive than their very high SES peers. This is somewhat unexpected since, a priori, one might suppose very high SES students to be better informed about higher education and thus not be much influenced by older schoolmates' higher education trajectory. Conversely, low SES students tend to be less aware of the higher education landscape ([Hoxby and Turner, 2013b](#)) and thus one could expect they would be more influenced by peers.

Third, the magnitude of spillovers vary across some characteristics of high schools. All high school tracks display spillovers of roughly the same magnitude, with slightly larger ones (in percentage terms) for the literature track. This result is quite noteworthy because it suggests that the acquisition of information about higher education choices is relevant

in very different contexts. That being said, spillovers are largest in small high schools (less than 30 students), and are decreasing in high school size. This could be the result of several explanations. Small high schools may exhibit closer relationships between students and their teachers, and thus teachers might be better aware of their students' higher education choices. As such they may encourage their subsequent classes to apply to the same college-majors as their past students. Another explanation could be that smaller high schools maintain better links with their alumni through, for example, annual alumni gatherings. A last explanation could simply be that smaller high schools are located in more rural areas where information about college-majors may be more scarce and therefore informational shocks are amplified to a much larger extent than in information-rich high schools located in large cities. Moreover, we find that high schools in the second and top quintile in the high school academic level distribution (measured as the median of its students' end of high school exam grades) exhibit the largest spillovers on applications. It is not clear exactly what could explain this result.

Fourth, we explore how spillovers differ across college-major characteristics. We find that spillovers are largest for public university, technical and vocational programs, but not for preparatory classes, which tend to be quite prestigious, nor for other types of college-majors. In line with these results, college-majors in the bottom 10% of selectivity (proxied by the median end-of-high school exam grade of enrolled students) exhibit the largest spillovers, and they are decreasing in selectivity. There are no spillovers for college-majors in the top 10%. This is somewhat surprising, as one may expect very selective college-majors to be those where some students may not dare to apply. This leads us to infer that students are learning about college-majors that they had been unaware of before rather than increasing their confidence to apply to prestigious college-majors.

Lastly, we assess how the interaction between high schools' and college-majors' characteristics may shape within-high school spillovers. The results suggest geographically close (less than 25 km) and moderately far (between 50 and 100 km) programs induce the largest spillovers. In terms of the intensive margin of applications and enrollment these moderately far programs display significant cross-cohort spillovers. Additionally, we find that students in low-achieving (bottom 25%) high schools are significantly more likely to follow an older schoolmate marginally admitted to a college-major in the top 25% of selectivity. This effect could be interpreted as raising the aspirations or awareness of these high schools' top performing students. Conversely, we find intriguingly that spillovers are quite large for top performing (top 25%) high schools for whom the marginally admitted student went to a college-major in the bottom quartile of selectivity.

In the last section of the paper we explore two mechanisms that may underpin our within-high school spillovers: (i) the role of teachers, and (ii) student role model effects. These mechanisms are not mutually exclusive but they lead to drastically different policy recommendations. We estimate the extent to which cross-cohort spillovers may be driven by teachers, for example, by recommending their past students' college-majors. Since we do not directly observe all of students' teachers, we test this mechanism in two complementary ways. First, we examine whether students are more likely to follow an older schoolmate if they share the same "*principal*" teacher. In France, each class is assigned a principal teacher who is in charge of the class' administrative duties over the course of the academic year, and in particular, helps and supervises students' higher education applications. Second, we assess whether students are more likely to follow an older schoolmate if they are in the same class identifier (e.g., senior class A, senior class B), an imperfect proxy for sharing the same set of teachers as that older schoolmate. We find that students sharing the same principal teacher or the same class identifier are equally likely to follow the marginally admitted older schoolmate's higher education choices as students with different teachers or principal teacher. This appears to suggest a rather limited direct role for teachers, at least in explaining the within-high school spillovers we document. This could be because teachers help their students by recommending a wide range of college-majors rather than only those of their past students.

Second, we attempt to disentangle whether our spillovers are more likely due to informational shocks or to role models effects. To test this, we assess whether the effects are larger for students sharing the same gender or socio-economic status as the marginally admitted older schoolmate. We interpret this test as capturing a role model effect rather than an information effect since, a priori, the marginally admitted student's gender or SES does not affect the informational content of his or her higher education trajectory but affects the way this information is perceived. We find strong evidence in favor of role model effects. Girls are significantly more likely to apply to college programs when the marginally admitted older schoolmate was a girl (+9%) but not when it was a boy (+3%, insignificant), while boys are more likely to follow a boy (+8%) but not a girl (+2%, insignificant). Similarly, low SES students are significantly more likely to apply to a degree when the marginally admitted older schoolmate was also of low SES background (+13%), but not when the latter is from a very high SES background (+1%, insignificant). However, very high SES students are largely unresponsive regardless of the SES of the treated older schoolmate. This is consistent with them being more knowledgeable about or having stronger preferences for college-majors.

Chapter 3: High-Achieving, Low-Income Students and Higher Education Financial Aid

Graduating from higher education provides one of the highest returns on investment an individual can make, especially when attending a selective institution (Bleemer, 2021; Black et al., 2023; Chetty et al., 2023). Yet, high-achieving, low-income students enrol at lower rates than their high-income peers, and when they do, they tend to attend lower quality institutions (Hoxby and Avery, 2013; Crawford et al., 2016; Dynarski et al., 2021; Hakimov et al., 2022; Campbell et al., 2022). This *undermatching* leads to large efficiency losses which could potentially be remediated by policy. Understanding the factors underlying these gaps is therefore crucial to design effective policy responses. Are high-achieving, low-income students less aware of the benefits of attending higher education, and specifically selective institutions? Do they lack information about relevant programs or simply do not have the self-confidence to apply? Or is it that they require additional financial resources to attend these selective colleges? If the former reasons prevail, then informational/motivational interventions should be favored. However, if financial constraints are the dominant explanation, then targeted financial support would be the preferred policy.

In this paper, I analyse whether additional financial aid can serve as an effective way of inducing high-achieving, low-income students to pursue higher education and enrol in high-quality institutions, as well as persist and graduate in a timely manner. Specifically, I estimate the effects of a national financial aid scheme, the *aide au mérite*, introduced in 2008 in France, which automatically granted an additional 1,800 euros annually, for 3 years at most (the duration of a bachelor's degree), to eligible students who enrolled in a higher education institution. The only criteria to be eligible to the *aide au mérite* were that the student (i) be eligible to the national need-based grant program, and (ii) score at least 16 out of 20 (i.e. in the top 4.7% of exam takers) at the French end of high school exam, the *Baccalauréat* (henceforth Bac).

The targeted population of students thus corresponds very closely to (Hoxby and Avery, 2013)'s definition of high-achieving, low-income students (top 4% of U.S. high school students, and in bottom parental income quartile). By design, the *aide au mérite* was awarded on top of need-based grants which included a tuition fee waiver and annual cash allowances up to 5,500 euros for the most disadvantaged students. As such the *aide au mérite* represented at least a 40% top up in monthly allowances, a sizable increase in financial support.

Using administrative data on the universe of students obtaining the Bac between 2009 and 2014, I exploit the sharp discontinuity in eligibility to the aide au mérite at the 16/20 Bac grade threshold in a regression discontinuity design. This enables me to estimate the causal effect of eligibility to this additional financial aid in the Bac year on enrollment, degree quality, persistence, graduation and academic performance in higher education as well as geographic mobility.

I find that being eligible to the aide au mérite in the Bac year had precisely estimated zero effects on enrollment, persistence or graduation from higher education. For most outcomes, I can reject effects as small as one to three percentage points. In this context, the enrollment margin is not particularly informative since, conditional on being eligible to a need-based grant, the enrollment rate around the 16 threshold is 94%. Moreover, students only become aware of their eligibility to the aide au mérite in July when Bac grades are released, which may limit the potential impact on enrollment. However, as in the U.S., persistence in higher education is a major concern in France. Around the 16 threshold, less than three out of four need-based grant eligible students are enrolled on time in 2nd year, and only just over half enrol in 3rd year on time. Thus, the null effects on persistence and graduation cannot be explained by students' late awareness of eligibility.

Additionally, I find no evidence that eligibility to the additional financial aid had an effect on the type or quality (proxied by the median Bac grade of students contemporaneously enrolling in the degree) of degree pursued. The null effects on degree quality remain for the degree enrolled in one year and two years later. This result rules out the hypothesis that eligible students become aware of the merit aid too late in the initial enrollment process but once aware subsequently choose to change tracks towards more selective degrees located in more expensive cities.

There is no discernible impact on other measures of higher education involvement such as the number of years enrolled in higher education or the highest level of study attained, nor on proxies for academic performance such as the likelihood of enrolling in a selective masters degree or the quality of the masters degree (again proxied by the median Bac grade of contemporaneous peers in the degree). Though I cannot observe students' undergraduate grades directly, this is indicative that academic performance does not appear to have been much influenced by eligibility to the aide au mérite. There is no clear sign of heterogeneous effects by gender or socio-economic background, suggesting these findings reflect true null effects and not heterogeneous effects that average out. This implies that high-achieving students' trajectories in higher education, even when they come from disadvantaged backgrounds, seem to be largely unaffected by the amount of finan-

cial support they receive. I do find evidence of positive effects on geographic location (Paris, and largest French cities) though the magnitude of the estimates are sensitive to the chosen bandwidth.

I exploit heterogeneity across specific subgroups to investigate three potential mechanisms that may underlie these null effects: (i) lack of information about eligibility to the aid, (ii) crowding out of parents' financial assistance, and (iii) the aid being awarded on top of need-based grants.

First, I find no evidence that the non-effect on *enrollment* might be driven by students being unaware of the policy. Since the aid was automatically granted to eligible students (conditional on enrolling in higher education), take up is not a concern. However, the aide au mérite was introduced at the same time as a vast reform of the need-based grants system and therefore may not have been as salient to students as this latter change. Yet, the estimates are not larger for more recent Bac cohorts who are very likely to have been more aware of the policy, nor are they are larger for students with more eligible high-school peers. This suggests that information deficits about the aide au mérite are unlikely to explain the null effect found on enrollment, though as discussed previously it could potentially be explained by students only becoming aware of eligibility late in the process. Since eligible students receive the financial aid once enrolled, there is no informational concerns for outcomes other than initial enrollment.

Second, I estimate the effects for students from the lowest-income families, who receive the highest need-based grants amounts but whose families are able to give them less than the amount of the aide au mérite on average (Grobon and Wolff, 2022). Thus, for these students, even if parental assistance is fully crowded out by the aide au mérite, they would still on net be better off financially. Admittedly, this will not necessarily be the case students whose parents' give them more than 200 euros monthly, and for whom the aide au mérite could theoretically be fully compensated by crowding out. I find no effects for the lowest-income students, suggesting that the overall null effects are unlikely to be the result of crowding out of parental financial contributions fully compensating the amount received from the aide au mérite. I cannot rule out potential interactions between eligibility and parent income that may not go through the crowding out channel, though one could expect that if there were any effects for a subgroup of students they would most likely be for the most disadvantaged students.

Lastly, I observe no evidence that students who are eligible only to the tuition fee waiver and no cash allowance as part of their need-based grant exhibit greater behavioral re-

sponses to eligibility to the aide au mérite than students who are eligible to more generous monthly cash allowances as part of their need-based grants. These results hold even when restricting to students with very similar parent incomes, suggesting these differences are not simply the result of differential parent incomes. This implies that the null effects are likely not completely driven by the aide au mérite being awarded on top of other financial aid, thus limiting its potential ability to have any effect.

This mechanism analysis indicates that the most likely explanation for the lack of observed effects is that high-achieving, low-income students are not marginal students, in the sense that their higher education outcomes are not contingent on the amount of financial aid they are eligible to. This is in line with a number of studies who consistently find that the impact of financial aid on higher education outcomes tends to be small (or null) for the highest ability students while effects for lower ability students are sizable (Goodman, 2008; Cohodes and Goodman, 2014; Fack and Grenet, 2015; Bettinger et al., 2019; Angrist et al., 2022). These findings highlight potential complementarities between financial aid and academic ability. A fruitful future research avenue would be to investigate more precisely how the effects of financial aid vary along the student ability distribution.

References

- Acciari, Paolo, Alberto Polo, and Giovanni L. Violante**, ““And Yet, It Moves”: Intergenerational Mobility in Italy,” *American Economic Journal: Applied Economics*, July 2022, 14 (3), 118–163.
- Aguirre, Josefa and Juan Matta**, “Walking in Your Footsteps: Sibling Spillovers in Higher Education Choices,” *Economics of Education Review*, February 2021, 80, 102062.
- Altmejd, Adam**, “Inheritance of Fields of Study,” 2023.
- , **Andrés Barrios-Fernández, Marin Drlje, Joshua Goodman, Michael Hurwitz, Dejan Kovac, Christine Mulhern, Christopher Neilson, and Jonathan Smith**, “O Brother, Where Start Thou? Sibling Spillovers on College and Major Choice in Four Countries,” *The Quarterly Journal of Economics*, June 2021, 136 (3), 1831–1886.
- Altonji, Joseph G., Erica Blom, and Costas Meghir**, “Heterogeneity in Human Capital Investments: High School Curriculum, College Major, and Careers,” *Annual Review of Economics*, September 2012, 4 (1), 185–223.
- Angrist, Joshua, David Autor, and Amanda Pallais**, “Marginal Effects of Merit Aid for Low-Income Students,” *The Quarterly Journal of Economics*, April 2022, 137 (2), 1039–1090.
- Aucejo, Esteban M., Claudia Hupkau, and Jenifer Ruiz-Valenzuela**, “Where versus What: College Value-Added and Returns to Field of Study in Further Education,” *Journal of Human Resources*, October 2022, pp. 0620–10978R1.
- Barrios-Fernández, Andrés**, “Neighbors’ Effects on University Enrollment,” *American Economic Journal: Applied Economics*, July 2022, 14 (3), 30–60.
- Bechichi, Nagui, Julien Grenet, and Georgia Thebault**, “D’Admission Post-Bac à Parcoursup : Quels Effets Sur La Répartition Des Néo-Bacheliers Dans Les Formations d’enseignement Supérieur ?,” Insee Références Édition 2021, INSEE, Paris 2021.
- Bettinger, Eric, Oded Gurantz, Laura Kawano, Bruce Sacerdote, and Michael Stevens**, “The Long-Run Impacts of Financial Aid: Evidence from California’s Cal Grant,” *American Economic Journal: Economic Policy*, 2019, 11 (1), 64–94.
- Björklund, Anders and Markus Jäntti**, “Intergenerational Income Mobility in Sweden Compared to the United States,” *American Economic Review*, 1997, 87 (5), 1009–1018.
- Black, Sandra E. and Paul J. Devereux**, “Recent Developments in Intergenerational Mobility,” in “Handbook of Labor Economics,” Vol. 4, Elsevier, 2011, pp. 1487–1541.
- , **Jeffrey T. Denning, and Jesse Rothstein**, “Winners and Losers? The Effect of Gaining and Losing Access to Selective Colleges on Education and Labor Market Outcomes,” *American Economic Journal: Applied Economics*, January 2023, 15 (1), 26–67.

- Bleemer, Zachary**, “Affirmative Action, Mismatch, and Economic Mobility after California’s Proposition 209,” *The Quarterly Journal of Economics*, December 2021, 137 (1), 115–160.
- Bowles, Samuel**, “Understanding Unequal Economic Opportunity,” *The American Economic Review*, 1973, 63 (2), 346–356.
- Campbell, Stuart, Lindsey Macmillan, Richard Murphy, and Gill Wyness**, “Matching in the Dark? Inequalities in Student to Degree Match,” *Journal of Labor Economics*, October 2022, 40 (4), 807–850.
- , –, –, and –, “Matching in the Dark? Inequalities in Student to Degree Match,” *Journal of Labor Economics*, 2023.
- Carrell, Scott and Bruce Sacerdote**, “Why Do College-Going Interventions Work?,” *American Economic Journal: Applied Economics*, July 2017, 9 (3), 124–151.
- Chetty, Raj, David J Deming, and John N Friedman**, “Diversifying Society’s Leaders? The Determinants and Causal Effects of Admission to Highly Selective Private Colleges,” NBER Working Paper 31492, National Bureau of Economic Research 2023.
- , **Nathaniel Hendren, Patrick Kline, and Emmanuel Saez**, “Where Is the Land of Opportunity? The Geography of Intergenerational Mobility in the United States,” *The Quarterly Journal of Economics*, 2014, 129 (4), 1553–1623.
- Cohodes, Sarah R. and Joshua S. Goodman**, “Merit Aid, College Quality, and College Completion: Massachusetts’ Adams Scholarship as an In-Kind Subsidy,” *American Economic Journal: Applied Economics*, October 2014, 6 (4), 251–285.
- Corak, Miles**, “Income Inequality, Equality of Opportunity, and Intergenerational Mobility,” *Journal of Economic Perspectives*, 2013, 27 (3), 79–102.
- , “The Canadian Geography of Intergenerational Income Mobility,” *The Economic Journal*, 2020, 130 (631), 2134–2174.
- Cortes-Orihuela, Javier, Juan Díaz, Pablo Gutierrez Cubillos, Pablo Troncoso, and Gabriel Villarroel**, “How Much Should We Trust TSTLS Intergenerational Mobility Estimates?: Evidence From A Developing Country,” *SSRN Electronic Journal*, 2022.
- Crawford, Claire, Paul Gregg, Lindsey Macmillan, Anna Vignoles, and Gill Wyness**, “Higher Education, Career Opportunities, and Intergenerational Inequality,” *Oxford Review of Economic Policy*, 2016, 32 (4), 553–575.
- Deutscher, Nathan and Bhashkar Mazumder**, “Intergenerational Mobility across Australia and the Stability of Regional Estimates,” *Labour Economics*, 2020, 66, 101861.
- and –, “Measuring Intergenerational Income Mobility: A Synthesis of Approaches,” *Journal of Economic Literature*, forthcoming.

- Dynarski, Susan, Cj Libassi, Katherine Micheltore, and Stephanie Owen**, “Closing the Gap: The Effect of Reducing Complexity and Uncertainty in College Pricing on the Choices of Low-Income Students,” *American Economic Review*, June 2021, 111 (6), 1721–1756.
- Estrada, Ricardo, Jeremie Gignoux, and Agustina Hatrick**, “Learning About Opportunity: Spillovers of Elite School Admissions in Peru,” 2022.
- Fack, Gabrielle and Julien Grenet**, “Improving College Access and Success for Low-Income Students: Evidence from a Large Need-Based Grant Program,” *American Economic Journal: Applied Economics*, April 2015, 7 (2), 1–34.
- Goodman, Joshua**, “Who Merits Financial Aid?: Massachusetts’ Adams Scholarship,” *Journal of Public Economics*, October 2008, 92 (10-11), 2121–2131.
- Grobon, Sébastien and François-Charles Wolff**, “Do Public Scholarships Crowd Out Parental Transfers ? Evidence from France,” *SSRN Electronic Journal*, 2022.
- Haider, Steven and Gary Solon**, “Life-Cycle Variation in the Association between Current and Lifetime Earnings,” *American Economic Review*, 2006, 96 (4), 1308–1320.
- Hakimov, Rustamdjan, Renke Schmacker, and Camille Terrier**, “Confidence and College Applications: Evidence from a Randomized Intervention,” 2022.
- Hastings, Justine, Christopher Neilson, and Seth Zimmerman**, “Are Some Degrees Worth More than Others? Evidence from College Admission Cutoffs in Chile,” Technical Report w19241, National Bureau of Economic Research, Cambridge, MA July 2013.
- Heidrich, Stefanie**, “Intergenerational Mobility in Sweden: A Regional Perspective,” *Journal of Population Economics*, 2017, 30 (4), 1241–1280.
- Hoxby, Caroline and Christopher Avery**, “The Missing “One-Offs”: The Hidden Supply of High-Achieving, Low-Income Students,” *Brookings Papers on Economic Activity*, 2013, (Spring).
- **and Sarah Turner**, “Expanding College Opportunities for High-Achieving, Low Income Students,” SIEPR Discussion Paper 12-014, Stanford Institute for Economic Policy Research, Stanford, CA 2013.
- Hoxby, Caroline M and Sarah Turner**, “Informing Students about Their College Options: A Proposal for Broadening the Expanding College Opportunities Project,” Technical Report 2013-03, The Hamilton Project 2013.
- Jacome, Elisa, Ilyana Kuziemko, and Suresh Naidu**, “Mobility for All: Representative Intergenerational Mobility Estimates over the 20th Century,” 2023.
- Jerrim, John, Álvaro Choi, and Rosa Simancas**, “Two-Sample Two-Stage Least Squares (TSTLS) Estimates of Earnings Mobility: How Consistent Are They?,” *Survey Research Methods*, 2016, 10 (2), 85–102.

Kirkeboen, Lars J., Edwin Leuven, and Magne Mogstad, “Field of Study, Earnings, and Self-Selection,” *The Quarterly Journal of Economics*, August 2016, 131 (3), 1057–1111.

Lefranc, Arnaud, “Intergenerational Earnings Persistence and Economic Inequality in the Long Run: Evidence from French Cohorts, 1931-75,” *Economica*, 2018, 85 (340), 808–845.

– **and Alain Trannoy**, “Intergenerational Earnings Mobility in France: Is France More Mobile than the US?,” *Annales d'Économie et de Statistique*, 2005, (78), 57–77.

Nybom, Martin and Jan Stuhler, “Biases in Standard Measures of Intergenerational Income Dependence,” *Journal of Human Resources*, 2017, 52 (3), 800–825.

1

Intergenerational Income Mobility in France: A Comparative and Geographic Analysis

This chapter is based on a paper co-authored with Louis Sirugue (PSE).

Abstract

We provide new estimates of intergenerational income mobility in France for children born in the 1970s using rich administrative data. Since parents' incomes are not observed, we employ a two-sample two-stage least squares estimation. We show, using the Panel Study of Income Dynamics, that this method slightly underestimates rank-based measures of intergenerational persistence. Our results suggest France is characterized by a strong persistence relative to other developed countries. 9.7% of children born to parents in the bottom 20% reach the top 20% in adulthood, four times less than children from the top 20%. We uncover substantial spatial variations in intergenerational mobility across departments, and a positive relationship between geographic mobility and intergenerational upward mobility. The expected income rank of individuals from the bottom of the parent income distribution who moved towards high-income departments is around the same as the expected income rank of individuals from the 75th percentile who stayed in their childhood department.

1. Introduction

TO what extent is the income of individuals related to that of their parents? This question has seen renewed interest both in the general public and in academia as rising income inequality raised concerns about equality of opportunity. Examining this link is essential to understand whether children from different socio-economic backgrounds are afforded the same opportunities. It also matters for economic efficiency, as high persistence across generations may reflect an inefficient allocation of talents (so-

called “Lost Einsteins”). Intergenerational persistence has now been estimated for a large number of countries, paving the way for insightful cross-country comparisons. Yet, much remains to be known for France, a country with relatively modest post-tax/transfers income inequality in international comparison and largely inexpensive higher education tuition fees.

The few existing studies for France only estimate the traditional intergenerational income elasticity (IGE), which captures the elasticity of child income with respect to parent income, and are based on small-sample surveys with self-reported incomes (Lefranc and Trannoy, 2005; Lefranc, 2018). Using a large sample combining census and tax returns data, we estimate two additional measures of intergenerational mobility: (i) the rank-rank correlation (RRC), increasingly prominent in the literature, which corresponds to the correlation between child and parent income percentile ranks, and (ii) transition matrices, which capture finer mobility patterns along the parent income distribution. While previous studies on France used self-reported labor earnings, we focus on household-level income measures. They provide a better depiction of one’s economic resources and allow the inclusion of children raised by single mothers. Integrating these improvements from the “new” intergenerational mobility literature enables us to conduct a detailed international comparison to rank France relative to other advanced economies for which comparable estimates are available.

In addition, we investigate the spatial variations in intergenerational mobility across the 96 metropolitan French departments. Such subnational analyses, pioneered by Chetty et al. (2014), help shed light on the mechanisms that may underlie income persistence across generations. Importantly, they highlight that national level estimates provide an incomplete assessment of a country’s intergenerational mobility. We make use of the panel dimension of our data to describe the geographic mobility patterns of individuals and study the relationship between geographic mobility and intergenerational mobility. We investigate the separate roles of moving to a higher-income department from that of climbing the income ladder within departments, conditional on parent income rank.

Our analysis is conducted on almost 65,000 children born between 1972 and 1981, and observed in the Permanent Demographic Sample (EDP). This rich administrative dataset allows us to implement the contributions discussed above and to convincingly address concerns related to lifecycle and attenuation bias (Haider and Solon, 2006; Black and Devereux, 2011; Nybom and Stuhler, 2017). Since parents’ incomes are not observed, we use a two-sample two-stage least squares (TSTSLS) estimation which consists in predicting parents’ incomes using other parents drawn from the same population but for whom income

is observed (Björklund and Jäntti, 1997). This method has been employed previously to estimate the IGE in the French context (Lefranc and Trannoy, 2005; Lefranc, 2018) as well as in many other countries (Jerrim et al., 2016, Table A1).

While studies typically use education and/or occupation to predict parent income, we make use of the richness of our data to also include detailed demographic characteristics of parents (French nationality dummy, country of birth, household structure, and birth cohort), and characteristics of the municipality of residence (unemployment rate, share of single mothers, share of foreigners, population, and population density). Our results are largely insensitive to the set of predictors. Parent income is then defined as the average¹ of father and mother predicted mean pretax wage over ages 35-45, and child income as pretax household income averaged over the same age range between 2010 and 2016. These two income definitions represent the most comprehensive household-level income definitions available for either generation.

TSTOLS Validation Exercise. Using the United States' Panel Study of Income Dynamics (PSID), we find that TSTOLS slightly underestimates rank-based measures of intergenerational persistence relative to what would be obtained if parent income were observed (OLS). The downward bias relative to the OLS estimate for the RRC ranges from 11% when education is the only predictor, to around 3-5% once occupation is also included. Subnational TSTOLS estimates are also fairly close to their OLS counterparts, though they tend to deviate more when the number of observations is small. Our results highlight that in settings like ours, where parent income cannot be directly observed, rank-based measures of intergenerational mobility obtained with TSTOLS likely provide lower bounds that are reasonably close to the true estimates. These findings confirm those obtained in different settings and samples by Cortes-Orihuela et al. (2022) and Jacome et al. (2023). We find that this reasoning also applies to the transition matrix.

National Results. Our main finding is that France exhibits relatively strong intergenerational income persistence compared to other developed countries. Our baseline estimate of the intergenerational elasticity in household income is 0.527, suggesting that on average, a 10% increase in parent income is associated with a 5.27% increase in child income. Put differently, if one's parents earn 10% more than the average of parents' incomes, then one is expected to preserve about 50% of that relative advantage. This estimate should be interpreted with caution considering our validation exercise suggests the TSTOLS IGE

¹See Section 3.3 for an explanation for why we take the *average* rather than the *sum*.

is significantly greater than the true estimate. Applying the correction factor we find, the IGE decreases to 0.396.

Moving to the rank-rank relationship, we find that the conditional expectation of child income percentile rank with respect to parent income percentile rank is linear throughout most of the parent income distribution, with steeper relationships at the tails. Our baseline estimate of the rank-rank correlation is 0.303, implying that a 10 percentile increase in parent income rank is associated, on average, with a 3.03 percentile increase in child income rank. This estimate is of similar magnitude to that found for Italy (0.3; [Acciari et al. \(2022\)](#)), somewhat smaller than for the United States (0.341; [Chetty et al. \(2014\)](#)), and markedly greater than existing estimates for other advanced economies such as Sweden (0.197; [Heidrich \(2017\)](#)), Australia (0.215; [Deutscher and Mazumder \(2020\)](#)) or Canada (0.242; [Corak \(2020\)](#)). Applying the correction factor we find in the validation exercise gives an RRC of 0.314 which does not affect France's relative position.

Intergenerational persistence, as captured by the transition matrix, is strongest at the tails of the parent income distribution: 9.7% of children from the bottom 20% of the parent income distribution reach the top 20% as adults. This probability is almost 4 times greater for children born to parents in the top 20% (38.4%). In comparison, the probability for a child born to a family in the bottom 20% to reach the top 20% in adulthood is 7.5% in the United States ([Chetty et al., 2014](#)) and 12.3% in Australia ([Deutscher and Mazumder, 2020](#)). Moreover, persistence at the top becomes stronger and stronger as we zoom in on the right tail of the parent income distribution. As with the RRC, the validation exercise suggests these estimates represent upper (lower) bounds on mobility (persistence).

We show that our baseline results are robust to potential biases. Foremost, we evaluate how sensitive they are to the parent income prediction specification. In particular, we check whether varying the set of predictors or using non-parametric estimation methods influences our estimates. IGE estimates are overinflated when using only education as a predictor, while the RRC and transition matrices remain surprisingly stable regardless of the set of predictors used. Slightly improved prediction from using flexible models does not quantitatively alter our estimates. Moreover, we assess our estimates' sensitivity to the lifecycle and attenuation biases by varying the ages at which child and parent incomes are measured as well as the number of parent income observations used. Our baseline results do not appear to under- nor over-estimate intergenerational mobility due to measuring child and/or parent incomes too early or too late in the lifecycle nor because of averaging incomes over too few years.

Subnational Results. We uncover substantial spatial variations in intergenerational mobility across departments, comparable to those observed across countries. We define individuals' location as their department of residence in 1990, when they are between 9 and 18 years old. Higher levels of mobility are typically found in the West of France, and lower levels in the North and South. While the IGEs range from 0.30 to 0.45 in departments in Brittany (West), they range from 0.42 to 0.70 in departments in Hauts-de-France (North). The distribution of department-level RRCs is tighter than that of IGEs, but displays very similar spatial patterns.

We also characterize departments' absolute upward mobility (AUM), defined as the expected income rank of children born to parents at the 25th percentile, which is obtained from the fitted values of the department-level rank-rank regression (Chetty et al., 2014). Absolute upward mobility ranges from the 36.8 in Pas-de-Calais (North) to 54.4 in Haute-Savoie (East). The Paris department stands out in terms of AUM (49.8) but exhibits around average intergenerational persistence levels in terms of IGE (0.51) and RRC (0.28). The cross-department correlation between the IGE and RRC is only 0.65, and -0.55 with AUM. This highlights the importance of using a variety of intergenerational mobility measures to characterize a country's income persistence across generations (Deutscher and Mazumder, forthcoming).

As a first step to understand the sources underlying these cross-department variations in intergenerational mobility, we undertake a simple correlational analysis. We find that absolute upward mobility exhibits much stronger relationships with department characteristics in general, than either the IGE or the RRC. This suggests that factors that affect absolute mobility might differ from those that affect relative mobility. The only characteristic consistently negatively correlated with intergenerational mobility is the unemployment rate. Intriguingly, we find no evidence of a within France "Great Gatsby Curve"² with respect to the IGE nor the RRC. This contrasts with findings from other countries (Acciari et al., 2022; Chetty et al., 2014; Corak, 2020).

Lastly, we conduct a descriptive analysis of the relationship between intergenerational income mobility and geographic mobility. We document important gains in expected income rank for movers, which are slightly decreasing in parent income rank. For children from families in the bottom decile, movers have an expected rank approximately 5.6 percentiles greater than stayers, while this difference is of roughly 4.4 percentiles for children

²The "Great Gatsby Curve" refers to the positive correlation between intergenerational income persistence (defined by the IGE) and income inequality (defined by the Gini index) found across countries (Corak, 2013).

from families in the top decile. These gains are partly attributable to movers locating in higher-income departments in adulthood relative to stayers, but also to movers reaching local ranks in their adulthood department that are further away from the rank of their parents in the childhood department. Destination departments are on average characterized by higher income levels than origin departments only at the tails of the parent income distribution. However, regardless of parent income rank, the absolute upward mobility gains associated with moving to a higher-income department appear to be large and increasing with average income in the destination department. All these findings combine self-selection and causal effects, and we leave the disentangling of these two channels for future research.

The rest of the article is organized as follows. Section 2 describes the intergenerational income mobility measures we estimate and the main sources of bias they are subject to. The data, the parent income prediction procedure and validation exercise, and the sample and variable definitions are presented in Section 3. Section 4 reports our baseline estimates at the national level, while Section 5 assesses their robustness to various sources of bias. In Section 6, we investigate the spatial variations in intergenerational income mobility, their correlation with local characteristics, and describe the relationship between geographic and intergenerational mobility. Section 7 concludes.

2. Measuring Intergenerational Mobility

Intergenerational income mobility can be characterized using a variety of statistics.³ In this section we (i) describe the statistics we employ, and (ii) discuss the two major biases inherent to most intergenerational persistence estimators, namely lifecycle bias and attenuation bias.

2.1. *Main Measures*

Intergenerational persistence measures primarily aim to characterize the joint distribution of children and their parents' lifetime incomes with a parsimonious set of practical statistics. We summarize intergenerational persistence using the following statistics.

³See for example [Corak \(2020\)](#), where nine statistics of intergenerational mobility are put into perspective. More elaborate discussions on the properties of the different intergenerational mobility estimators can also be found in [Black and Devereux \(2011\)](#), [Chetty et al. \(2014\)](#), [Nybom and Stuhler \(2017\)](#), and [Deutscher and Mazumder \(forthcoming\)](#).

Intergenerational Income Elasticity (IGE). The traditional intergenerational income elasticity is obtained by regressing children’s log lifetime income on their parents’ log lifetime income. An IGE of 0.4 implies that a 10% increase in parent income is associated, on average, with a 4% increase in child income. Importantly, this estimator is sensitive to differences in inequality across generations. This can be seen in the following equation, where y_p and y_c are parent and child log lifetime incomes:

$$\text{IGE} = \frac{\text{Cov}(y_c, y_p)}{\text{Var}(y_p)} = \text{Corr}(y_c, y_p) \times \frac{\text{SD}(y_c)}{\text{SD}(y_p)}. \quad (1.1)$$

The empirical literature has highlighted that IGEs are particularly sensitive to lifecycle and attenuation biases, sample selection criteria, non-linearities along the parent income distribution, income definitions, and to the treatment of negative/zero incomes (Couch and Lillard, 1998; Chetty et al., 2014; Landersø and Heckman, 2017; Helsø, 2021).

Rank-Rank Correlation (RRC). The increasingly popular rank-rank correlation is obtained by regressing children’s percentile rank in lifetime income on their parents’ percentile rank in lifetime income. A RRC of 0.4 means that a 10 percentile increase in parent rank is associated, on average, with a 4 percentile increase in child rank. Unlike the IGE, the RRC is unaffected by inequality levels in either generation. This can be seen in the following equation, where p_p and p_c are parent and child percentile ranks in their respective lifetime income distributions:

$$\text{RRC} = \frac{\text{Cov}(p_c, p_p)}{\text{Var}(p_p)} = \text{Corr}(p_c, p_p) \times \frac{\text{SD}(p_c)}{\text{SD}(p_p)} = \text{Corr}(p_c, p_p). \quad (1.2)$$

Consequently, the greater the degree of inequality in the child generation relative to the parent generation, the greater the IGE relative to the RRC. In addition, the same RRC in two countries with large differences in inequality would hide that in one country the distance between ranks in monetary terms is actually much larger than in the other. The RRC owes its recent popularity to its robustness to specification variations, common biases, and treatment of negative/zero incomes (Dahl and DeLeire, 2008; Chetty et al., 2014; Nybom and Stuhler, 2017).

Transition Matrices. To get a finer picture, one can use transition matrices, which report the probability of ending up in a given quantile as an adult conditional on coming from

a family in a given quantile. Typically, they are reported by quintile and are of particular interest to seize non-linearities in children mobility along the parent income distribution.

2.2. *Main Sources of Bias*

The vast majority of currently available data sources do not cover the whole lifetime of children's and/or parents' incomes, leading researchers to approximate lifetime income based on shorter time spans. This data limitation generates the following two fundamental biases, which we extensively investigate in Section 5.

Attenuation Bias. A direct implication of relying on a limited number of income observations to approximate parent lifetime income is the attenuation bias arising from classical measurement error (Solon, 1992; Zimmerman, 1992). This leads to downward-biased estimates of intergenerational persistence. Mazumder (2005, 2016) and Nybom and Stuhler (2017) find that the attenuation bias can be very large for the IGE but affects the RRC only mildly, while O'Neill et al. (2007) show that it affects most the corner elements of the transition matrix. The common solution to lessen this bias is to average parent income over as many years as possible.

Lifecycle Bias. The second common bias relates to the age at which child and parent incomes are observed (Grawe, 2006; Haider and Solon, 2006). In particular, lifecycle bias arises in the presence of heterogeneous age-income profiles, which is observed empirically as high lifetime income individuals tend to experience steeper earnings profiles than low lifetime income individuals. As such, observing child or parent incomes either too early or too late in the lifetime is likely to bias intergenerational persistence estimates. The IGE is particularly sensitive to lifecycle bias, especially if incomes are measured before age 35, while it affects the RRC only moderately so long as incomes are measured at least in the late 20s/early 30s. Just as for the attenuation bias, the corner elements of the transition matrix are most sensitive to lifecycle bias (Chetty et al., 2014; Nybom and Stuhler, 2016, 2017).

3. Data

We use data from the Permanent Demographic Sample (EDP), which combines several administrative data sources on individuals born on the first four days of October. We

refer to individuals born on one of these days as *EDP individuals*. We describe below the most relevant details for each data source we use and provide additional technicalities in Appendix A.

Civil Registers. They contain information from birth certificates of EDP individuals and their children, including gender, date and place of birth, and parents' date and place of birth, nationality and occupation.

1990 Census. It contains socio-demographic information about EDP individuals and members of their household. Importantly, it reports parents' education level, occupation, and other demographic characteristics if EDP individuals live with their parents in 1990.

All Employee Panel. It gathers worker-year level information on all private (since 1967) and public (since 1988) sector employees in metropolitan France, except those in the agricultural sector. Prior to 2001, only individuals born on an even year are covered. Our results are robust to the late coverage of civil servants (see Appendix C.1).

Tax Returns. They provide tax information on incomes earned between 2010 and 2016 for individuals in dwellings where an EDP individual is known either from their income tax form or their main housing tax. Income variables are available both at the household level and at the individual level. An advantage of the information being gathered at the dwelling level is that household income is observed for all couples, regardless of whether they file their taxes jointly.

3.1. *Parent Income Prediction*

The measures of intergenerational mobility laid out in Section 2.1 cannot be estimated directly with our data since we do not observe parents' incomes. We therefore rely on the two-sample two-stage least squares (TSTLS) strategy introduced by Björklund and Jäntti (1997), and previously used in the French context by Lefranc and Trannoy (2005) and Lefranc (2018), and in many other countries (Jerrim et al., 2016, Table A1). It consists in predicting individuals' parents' incomes from a sample of other parents whose incomes are observed using a set of common observed characteristics. We refer to these other parents as *synthetic* parents.

Let Z denote a set of characteristics observed both for parents and synthetic parents. Their log lifetime incomes y can be expressed as:

$$y_i = \beta Z_i + \varepsilon_i. \quad (1.3)$$

We estimate this first-stage equation by OLS on our sample of synthetic parents, and predict parents' log lifetime incomes using the resulting $\hat{\beta}$ as $\hat{y}_i = \hat{\beta} Z_i$. Z includes parents' (i) education (8 categories), (ii) 2-digit occupation (42 cat.; includes inactivity status), (iii) demographic characteristics (birth cohort, French nationality dummy, country of birth (6 cat.), and household structure (6 cat.)), and (iv) characteristics of the municipality of residence (unemployment rate, share of single mothers, share of foreigners, population, and population density). For the geographic analysis, we drop the municipality characteristics to ensure they do not spuriously drive any spatial patterns, though this has virtually no impact on the estimates. All characteristics are observed in the 1990 census. To reduce the potential for lifecycle and attenuation bias, synthetic parents' income is defined as average pretax wage between 35 and 45 with at least 2 income observations over this age range in the All Employee Panel. The model is estimated separately on synthetic mothers (adj. $R^2 = 0.37$) and fathers (adj. $R^2 = 0.36$). We extensively test the robustness of our baseline results to using more flexible models and to varying the set of first-stage regressors in Section 5.1.

Method Validity. To assess how reliable TSTSLS estimates are relative to their OLS counterparts (i.e., using *observed* parent income), we need a dataset that includes parents' actual incomes as well as predictors of parents' incomes. Since such a dataset does not exist for France, we follow [Jerrim et al. \(2016\)](#), [Bloise et al. \(2021\)](#) and [Jacome et al. \(2023\)](#), and conduct a validation exercise using the United States' Panel Study of Income Dynamics.⁴ We describe this analysis in detail in Appendix B. We provide comparisons both at the national level and, due to sample size constraints, by Census Bureau regions (Northeast, Midwest, South, and West). Our sample and definition choices aim to be as close as possible to our main analysis setting while at the same time maximizing sample size.

Specifically, our sample of children consists in individuals born between 1963 and 1988. We define parent income as the sum of father and mother mean labor income over ages 30-50, and child income as mean family total income over ages 30-50. The results are

⁴[Acciari et al. \(2022\)](#) and [Cortes-Orihuela et al. \(2022\)](#) also conduct validation exercises of TSTSLS using administrative data from Italy and Chile respectively.


quantitatively similar when computing parent and child incomes over ages 35-45 as in the main analysis, despite the smaller sample size. We use education, 3-digit occupation (including inactivity status), birth year, race, and state of residence as first-stage predictors. These predictors are the closest we could find to those used in the main analysis.

Figure 1.1 presents the main results from our validation exercise. At the national level, the TSTSLS RRC estimate (0.459) is 4% smaller than the OLS estimate (0.476), a very moderate difference. Moreover, and importantly, the TSTSLS estimate of the RRC appears to understate persistence, i.e., they provide an upper bound for intergenerational mobility, as also found by Cortes-Orihuela et al. (2022) and Jacome et al. (2023). The same applies for estimates of the transition matrix presented in Appendix Figure B.3. At the Census Region level, the RRC obtained by TSTSLS are again reasonably similar to those obtained by OLS, with a slightly larger underestimation for the Northeast and West regions where the number of observations is smaller. The same applies to absolute upward mobility, defined as the expected rank of children from families at the 25th percentile.

The TSTSLS RRC estimate is smaller than the true OLS estimate likely because parents from the very top (bottom) of the income distribution can only be mispositioned downwards (upwards) when using predicted incomes. Assuming a monotonic relationship between parents and child income ranks, this mechanically flattens the rank-rank relationship and biases the rank-rank correlation downwards. This can be seen in Figure 1.2, which shows the conditional expectation of out-of-sample predicted labor income rank with respect to observed labor income rank, as well as the interquartile range of the prediction. Indeed, percentile ranks tend to be overestimated at the bottom of the parents income distribution and underestimated at the top. We obtain very similar out-of-sample predictions in the EDP as in the PSID, suggesting we can reasonably apply the estimated TSTSLS biases of our validation exercise to the main analysis. Note that the IGE is sensitive to another bias because, all else equal, it is decreasing in the variance of parents' incomes (as highlighted in equation (1.1)). As such, since the distribution of predicted parent incomes is narrower than the true distribution, this puts an upwards pressure on the IGE.

Inference. Since we are in a two-stage setting, standard inference is inappropriate. Inoue and Solon (2010) derive an analytical formula for TSTSLS standard errors. However, their method cannot be applied in our setting as we use a non-standard transformation of the first-stage outcome variables. Indeed, because labor income is observed for synthetic parents *individually* but is not observed for their spouse, we can only estimate equation


Figure 1.1. OLS vs. TSTSLs RRC - National and Census Regions in the United States



Notes: This figure presents rank-rank correlations obtained when parent income is observed (OLS) and when it is predicted using two-sample two-stage least squares (TSTSLs), at the national level and by Census Bureau Regions in the United States. They are computed on the Panel Study of Income Dynamics (PSID). The sample used is restricted to children born between 1963 and 1988 who are observed at least once as children in a family unit and at least once as a reference person or partner in a family unit over ages 30-50. Child income is the mean of family total income over ages 30-50. Parent income is the sum of father and mother (predicted) mean labor income over ages 30-50. For TSTSLs estimates, parent income is predicted separately for males and females using an OLS model including education (7 cat.; highest years of school completed), 3-digit occupation (334 cat.; most common occupation (incl. inactivity status) between 30 and 50 years old), parents' demographic characteristics in 1990 (birth cohort and race (5 cat.; most recent observation) and state fixed effects (most common state of residence between 30 and 50 years old). The fitted lines correspond to the regression line obtained on the microdata. We report coefficients and naive standard errors (in parenthesis) obtained from OLS regressions of child income rank on parent income rank with child cohort fixed effects, on the microdata for the full sample.

(1.3) on *individual* income. We then aggregate mother and father predicted incomes to obtain a measure of *household* income, which we use as the regressor in the second stage rather than using the fitted values from the first stage as is. We thus report bootstrapped standard errors for all individual-level regressions, which, for the same reason, cannot be clustered at the family level. Specifically, we draw one bootstrap sample for synthetic fathers and one for synthetic mothers separately. We then run the first-stage regression,

Figure 1.2. Out-of-Sample Predicted Labor Income Rank - France (EDP) and United States (PSID)



Notes: This figure presents the conditional expectation of out-of-sample predicted labor income rank with respect to observed labor income rank, for both the PSID validation exercise (United States - PSID) and our own parent income prediction (France - EDP). See Figure 1.1's notes for details on data, sample and income definitions for the PSID analysis, and Figure 1.3's note for details on our analysis (EDP).

and predict parent income on a bootstrap sample of children. We iterate this process 1,000 times. These bootstrapped standard errors are of the same order of magnitude though slightly larger than naive ones.

3.2. *Sample Definitions*

Hereinafter we rely on the Permanent Demographic Sample (EDP) to estimate intergenerational persistence in France. Our samples of interest are defined as follows.

Sample of Children. It consists of EDP individuals who are (i) born between 1972 and 1981 in metropolitan France,⁵ (ii) observed with their parents in the 1990 census, (iii) whose parents are neither farmers nor in a liberal profession⁶, and (iv) observed in the tax

⁵Metropolitan France refers to the part of France that is geographically in Europe.

⁶Liberal professions encompass activities that are not salaried, agricultural, commercial or artisanal, and carried out by self-employed service providers (e.g., lawyers, notaries, private doctors, etc.). 5.08% of EDP

returns data at least once between 35 and 45 years old.⁷ Restriction (i) is made to observe individuals with their parents in the 1990 census⁸ and to have a reasonably large sample size for the subnational analysis. Restriction (ii) enables us to retrieve their parents' characteristics, and (iii) is due to the fact that farmers and liberal professions are not covered by the All Employee Panel from which we obtain synthetic parent income. Restriction (iv) aims to minimize lifecycle bias. The final sample contains 64,571 children.⁹ Overall, they have very similar socio-economic characteristics as the representative sample of EDP individuals satisfying only restriction (i), except for under-representing children of farmers by definition, as shown in Appendix Section C.1.

Sample of Synthetic Parents. It is constructed such that synthetic parents come from the same overarching population as actual parents. It therefore consists of EDP individuals who (i) had at least one child born between 1972 and 1981 in metropolitan France, (ii) are observed in the 1990 census, (iii) are neither farmers nor in a liberal profession in 1990, and (iv) have at least two pretax wage observations between 35 and 45 years old in the All Employee Panel.¹⁰ As such our sample excludes individuals born in an odd year since they were not covered by the All Employee Panel prior to 2001. The final sample contains 31,423 synthetic parents.¹¹

Descriptive Statistics. Appendix Table F.6 provides statistics on our sample of synthetic parents and children. On average, fathers are around 42 in 1990 and mothers 39. This assures that we predict income based on observable characteristics measured sufficiently late in their lifecycle.

individuals satisfying (i) and (ii) have at least one parent who is a farmer and 2.41% have at least one parent who is in a liberal profession. As raised by Lefranc (2018), the fact that farmers tend to face relatively low incomes and a strong occupational inheritance (Lefranc et al., 2009) makes the exclusion of farmers likely to bias intergenerational persistence downwards.

⁷6.73% of EDP individuals satisfying (i) and (ii) are not observed in the tax returns data between 35 and 45 years old.

⁸See Appendix Figure E.1 for the position in the family in the 1990 census by child birth cohort.

⁹See Appendix Table F.1 for the sample size at each additional restriction. Parent income cannot be predicted for 23 children because one of their parents has an occupation not represented in the sample of synthetic parents of the corresponding gender, hence the very slight discrepancy with this table.

¹⁰In Appendix Table F.2 we compare average characteristics of parents and synthetic parents. To ensure appropriate comparability of the two samples, no restriction on wage observations for synthetic parents or children is applied. Average characteristics are remarkably similar for most variables, even for 2-digit occupation (Appendix Table F.3), which confirms the assumption that actual and synthetic parents are random subsets of the same population.

¹¹See Appendix Table F.4 for the sample size at each additional restriction.

3.3. Variable Definitions

The variables we use are constructed as follows. All incomes are expressed in 2015 euros, and are measured before taxes but after the deduction of employer- and employee-level payroll taxes.

Parent Income. We define the income of one parent as predicted average pretax wage over ages 35 to 45. This income is predicted according to the methodology described in Section 3.1. We then compute income at the household level (regardless of marital status) by taking the average of father and mother predicted incomes if the child is observed with both parents in the 1990 census, and income of the only parent otherwise. We take the *average* of father and mother predicted incomes rather than the *sum* (the standard in the literature), to correct for the fact that otherwise single-headed households would be over-represented in the bottom of the income distribution (when using the sum, there are virtually no single-headed households above rank 50). Indeed, while in other studies parent income is typically observed repeatedly over several years, in our setting a parent observed as single in 1990 can by definition only be predicted their *individual* income for their entire lifetime even if their marital status actually changes later on. We refer to this income definition as parent household wage and use it as our main parent income measure. We also report results using father predicted income, which we refer to as father wage.

Child Income. Our main measure of child income, computed from the tax returns, corresponds to the sum of labor earnings (wages and self-employment income), taxable and imputed non-taxable capital income¹², unemployment insurance, retirement, and alimony, at the household level.¹³ Just as for parents, a household is defined as individuals living in the same dwelling. To mitigate the potential for lifecycle bias, we average over 2010-2016 only for incomes declared when the individual is between 35 and 45 years old. We refer to this income definition as household income and use it as our main child income measure. We also report results using the following alternative child income defi-

¹²Financial incomes not subject to any tax reporting are predicted by the French National Institute of Statistics and Economic Studies (INSEE) from a model estimated on the *Enquête Patrimoine*. In particular, they predict capital income for seven financial products (various tax-exempt savings accounts and life insurance) using household-level observed characteristics (income, age, family situation, ...). Excluding this income source from our child income definition does not affect the results.

¹³Social benefits such as family allowances, social minima (e.g., RSA, disability benefits) and housing benefits are not included in our main measure of child income.

nitions: (i) household wage, which is equivalent to the parent household wage definition, (ii) individual income, which we define as the sum of all individual-level incomes: labor earnings (wages and self-employment income), unemployment benefits, retirement, and alimony, and (iii) individual wage.

Income Definition Discussion. Our preferred parent and child income definitions represent the most comprehensive household-level income definitions possible for either generation. Defining incomes at the household level is important in order to (i) better capture the economic conditions of individuals and their parents, (ii) allow the inclusion of children raised by single mothers, and (iii) enable the analysis of daughters, whose labor incomes alone may not be an appropriate measure of their economic outcomes. These income definitions are not identical but the results are qualitatively similar when using the same income definition, household wage, for both children and parents.

Percentile Ranks. We rank children within their birth cohort, and parents relative to other parents with children in the same birth cohort. To avoid individuals (in a given cohort) earning the same income (e.g., 0, or the minimum wage) being assigned different income ranks, we define the income rank of such individuals as the ceiling of the median income rank of individuals with that income level.¹⁴

4. Results at the National Level

We start by analyzing intergenerational mobility at the national level. For our baseline results, we use data on children born on the first four days of October between 1972 and 1981 and measure parent income as household-level predicted average annual pretax wage over ages 35-45, and child income as pretax household income averaged over the same age range between 2010 and 2016. We include child birth cohort fixed effects in the log-log and rank-rank regressions.¹⁵

¹⁴For example, if there are 3.65% of children with zero income, their median rank is 2, and thus they are assigned a rank of 2. In our samples, 0.06% of children have negative or zero household income (see Appendix Table F.5), while no parent has negative or zero predicted wage.

¹⁵In practice, these fixed effects have virtually no influence on the coefficients of interest.

4.1. *Intergenerational Income Elasticity (IGE)*

Figure 1.3 panel A displays the conditional expectation of log child income with respect to log parent income. Children with negative or zero incomes are excluded. This is of minor importance when defining child income as household income as such cases are exceedingly rare. Nonetheless, we assess the influence of zero incomes in Appendix Figure C.9. The log-log CEF is pretty linear throughout the middle 80% of the parent income distribution, with some mild non-linearities at the tails.¹⁶ This S-shaped relationship is also observed in the United States (e.g., Chetty et al. (2014)), Denmark (e.g., Helsø (2021)) or Sweden (e.g., Björklund et al. (2012)). It implies that the elasticity is not constant over the whole parent income distribution, with smaller magnitudes at the tails, and is sensitive to the inclusion or exclusion of parents at the tails of their income distribution.¹⁷

Our baseline IGE estimate is 0.527, meaning that a 10% increase in parent income is associated, on average, with a roughly 5% increase in child income. This estimate should be interpreted with caution as our validation exercise presented in Section 3.1 suggests TST-SLS estimates of the IGE can be quite inflated relative to the true value. Thus this baseline IGE is not well-suited for international comparisons. Appendix Figure E.2 shows our estimates of the intergenerational income elasticity for every child and parent income definition, and for sons and daughters separately. Our father-son wage IGE estimate is relatively similar to existing ones for France despite important differences in methodology and data (see Appendix Table F.7). Intergenerational persistence estimates are larger for household income than for individual income or wage, which could be the result of assortative mating. IGEs are very similar when defining parent income as father wage, despite the fact that by construction, estimates based on father wage exclude children only observed with their mother in the 1990 census (about 10% of observations). The IGE is significantly lower for sons (0.478) than for daughters (0.577). This phenomenon is not systematic across countries, but is also observed in Germany (Bratberg et al., 2017) and the Netherlands (Carmichael et al., 2020), for instance.

4.2. *Rank-Rank Correlation (RRC)*

Figure 1.3 panel B plots the conditional expectation of child income rank with respect to parent income rank. It is relatively linear, with slight non-linearities at the tails as observed in many countries (Chetty et al., 2014; Bratberg et al., 2017; Helsø, 2021).

¹⁶Appendix Figure C.2 shows that these non-linearities are not driven by the set of first-stage predictors.

¹⁷Appendix Figures C.10a and C.10c show how trimming the top and bottom of the parent/child income distribution influences our estimates.

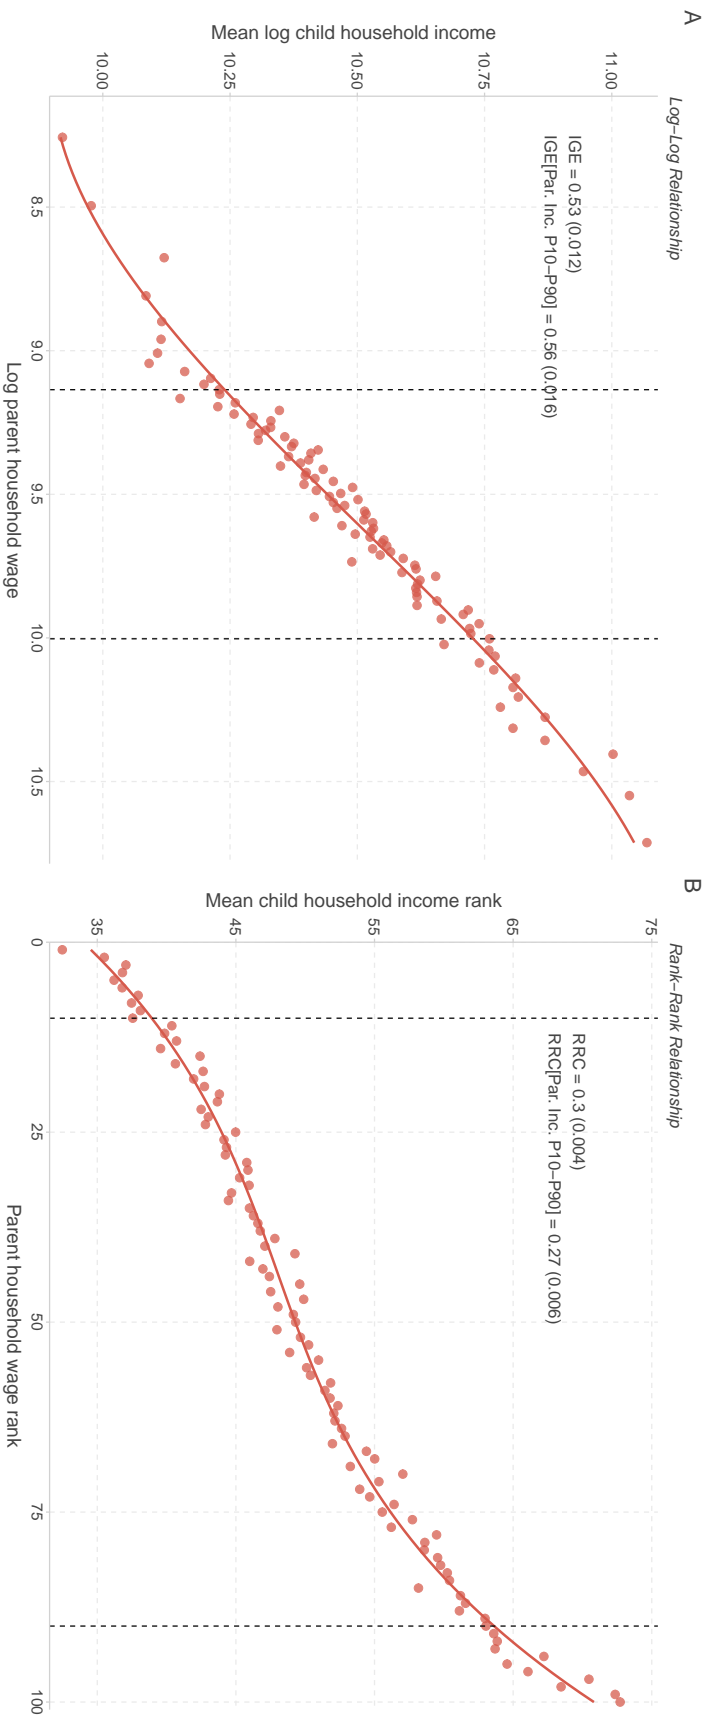


Figure 1.3. Conditional Expectation Functions for Log-Log and Rank-Rank Relationships in France

Notes: This figure presents non-parametric binned scatter plots of the relationship between log child income and log parent income (panel A), and child income rank and parent income rank (panel B) in France. It is computed on the Permanent Demographic Sample, a dataset of individuals born on the first four days of October. The sample used is restricted to children born between 1972 and 1981. Child income is the mean of 2010–2016 household income (with age restricted to 35–45). Parent income is the sum of each parent predicted wage divided by the number of parents. Parent income is predicted separately for males and females using an OLS model including parents' education (8 cat.), 2-digit occupation (42 cat.), demographic characteristics in 1990 (birth cohort, French nationality dummy, country of birth (6 categories), and household structure (6 cat.)) and characteristics of the municipality they lived in in 1990 (unemployment rate, share of single mothers, share of foreigners, population, and population density). These municipality characteristics are excluded for the geographic analysis. It is estimated on a sample of synthetic parents whose average wage at ages 35–45 (at least 2 income observations) is used as the dependent variable. Incomes are in 2015 euros. To construct panel A, children with negative or zero incomes are excluded (0.06% of the sample) and we bin parent incomes into 100 equal-sized bins and plot mean log child income versus mean log parent income within each bin. To construct panel B, children are ranked relative to other children in the same birth cohort while parents are ranked relative to other parents with children in the same birth cohort. We then plot mean child income rank versus parent income rank. The dashed lines represent the 10th and 90th percentiles of parents' income. We report coefficients and bootstrapped standard errors (in parenthesis) obtained from OLS regressions of log child income on log parent income (panel A) and child income rank on parent income rank (panel B), both with child cohort fixed effects, on the microdata for the full sample and for parents between the 10th and 90th percentiles. The fitted line is a 3rd order polynomial fit through the conditional expectations.

Our baseline estimate of the rank-rank correlation is 0.303, meaning that a 10 percentile increase in parent income rank is associated, on average, with a 3.03 percentile increase in child income rank. Note that this estimate corresponds to a lower bound, as the validation exercise suggests the TSTOLS methodology slightly underestimates the RRC. Applying the estimated correction factor of 3.7% leads to a corrected baseline RRC coefficient of 0.314. Appendix Figure E.3 shows our baseline estimates of the rank-rank correlation for every child and parent income definition, and for sons and daughters separately. The estimates are slightly higher for daughters (0.310) than for sons (0.296), and are also slightly higher when defining parent income as household wage rather than as father wage. The estimates are smaller when defining child income as household wage or individual income and smallest when using individual wage, a pattern observed in other countries (Chetty et al., 2014; Deutscher and Mazumder, 2020; Landersø and Heckman, 2017), again possibly due to assortative mating.

To the best of our knowledge, this is the first time the RRC is estimated for France.¹⁸ In Table 1.1 we compare RRC estimates for countries for which estimates exist (see Appendix Figure E.5 for a visual representation). To enable comparability we only keep studies which pool sons and daughters together, define parent income at the household level and use comprehensive income definitions. Note that for child income some studies only observe *individual* rather than *household* income which might result in lower RRC estimates (as we find for France, and Chetty et al. (2014) for the United States). Even though they are not directly comparable due to important differences in data and sample selection rules, we believe that it is a relevant exercise given the relative stability of the RRC to specification variations and common data limitations.

This international comparison suggests that (i) France exhibits strong persistence across generations in international comparison, given that it is the country with the second highest available RRC estimate behind the United States, and (ii) there is less variation across countries in the rank-rank slope than in the intergenerational elasticity, which is coherent with the fact that the RRC is not influenced by changes in inequality across generations, and is less sensitive to sample restrictions.

¹⁸A recent report (in French) by Abbas and Sicsic (2022) now also provides rank-based intergenerational mobility estimates for France. They use the same data as us and their sample consists in individuals born in 1990 (i) who are still claimed as dependent in their parents' tax return at age 20, (ii) whose parents' income can be observed around age 50, and (iii) whose individual income is observed at age 28 in their own tax return. They compare their results to ours and despite different sample definitions, when using the same income definition and measuring child income at the same age (i.e., 28), they find very similar results.

Country	RRC ↓	# obs.	Data	Child Income Definition ¹	Child Cohort	Child Age or Year at Income Measurement	Parent Age or Year at Income Measurement	Source
Switzerland	0.14	667,047	Social Security Records	Average total pretax <i>individual</i> income	1967-1982	32-34	when child between 15-20	Kalambaden and Martinez (2021, Table 3)
Switzerland	0.14	923,262	Social Security Records	Average total pretax <i>individual</i> income	1967-1984	30-33	when child between 15-20	Chuard-Keller and Grassi (2021, Figure 1)
Spain	0.195	1,492,107	Atlas de Oportunidades	Total pretax <i>individual</i> income	1980-1986	2016	1998	Soria Espin (2022, Figure 1)
Sweden	0.197	778,484	SIMSAM database ²	Average total pretax <i>individual</i> income	1968-1976	32-34	34-50	Heidrich (2017, Table 2)
Denmark	0.203	157,543	Danish register data	Average total pretax <i>family</i> income	1980-1982	2011-2012	1996-2000	Helsø (2021, Table 1)
Australia	0.215	1,025,800	Federal income tax returns	Average total pretax <i>family</i> income	1978-1982	2011-2015	1991-2001	Deutscher and Mazumder (2020, Table 2)
Sweden	0.215	252,745	35% random sample from admin. data	Average total pretax <i>household</i> income	1957-1964	1996-2007 ³	1978-1980	Bratberg et al. (2017, Table 3)
Norway	0.223	324,870	Full population admin. data	Average pretax <i>family</i> earnings	1957-1964	1996-2006	1978-1980	Bratberg et al. (2017, Table 3)
Canada	0.242	2,115,150	Intergenerational Income Data	Average total pretax <i>family</i> income	1963-1970	2004-2008	when child between 15-19	Corak (2020, Table 5)
Germany	0.245	1,128	German Socio-Economic Panel	Average total pretax <i>household</i> income	1957-1976	2001-2012	1984-1986	Bratberg et al. (2017, Table 3)
Denmark	0.253	≈ 410,000	Danish register data	Average total pretax <i>individual</i> income	1973-1979	2010-2012	when child between 7-15	Landersø and Heckman (2017, Table A6)
Denmark	0.257	205,625	Full population admin. data	Average total pretax <i>individual</i> income	1973-1975	2010-2012	when child between 7-15	Eriksen (2018, Table 3.2)
Italy	0.30 ⁴	1,719,483	Electronic database of Personal Income returns	Average total pretax <i>individual</i> income	1979-1983	2016-2018	1998-2000	Acciari et al. (2022, p.145)
France	0.303⁵	64,571	Permanent Demographic Sample	Parents: (predicted) <i>household</i> wage; Children: average total pretax <i>household</i> income	1972-1981	2010-2016 (between 35-45)	35-45	
United States	0.341	9,867,736	Federal income tax records, 1996-2012	Average total pretax <i>family</i> income	1980-1982	2011-2012	1996-2000	Chetty et al. (2014, Table 1)
United States	0.395	6,414	NLSY79	Average total pretax <i>family</i> income (self-reported)	1957-1964	1996-2008	1978-1980	Bratberg et al. (2017, Table 3)

Notes:

¹ The parent income definition is always at the family level.

² Swedish Initiative for Research on Microdata in the Social and Medical Sciences.

³ Only even years.

⁴ This estimate corresponds to the one when adjusting for lifecycle bias, incomplete coverage of taxpayers and tax evasion as reported on p.28. The baseline RRC estimate reported in Table 3 is 0.22.

⁵ Assuming that the bias induced by the TSTIS methodology is the same in France as in the United States, our validation exercise performed on the PSID suggests the OLS counterpart to our baseline estimate would equal to $0.303 \times \frac{0.476}{0.289} = 0.314$ (see Figure 1.1).

Table 1.1: Rank-Rank Correlation in International Comparison

4.3. Transition Matrices

The last measure of intergenerational income persistence we estimate is a quintile-by-quintile transition matrix, which documents the conditional probabilities of being in each income quintile as an adult given any parent income quintile. Figure 1.4 presents our baseline estimates of the transition matrix for France, along with available estimates for the United States (Chetty et al., 2014) and Australia (Deutscher and Mazumder, 2020). To the best of our knowledge, this is the first time transition matrices are estimated for France.¹⁹




Figure 1.4. Baseline Quintile Transition Matrix for Different Countries

Notes: The first panel of this figure presents our baseline intergenerational transition matrix estimates. Bootstrapped standard errors are presented in Appendix Figure E.4. See Figure 1.3’s notes for details on data, sample and income definitions. Each cell documents the share of children belonging to the quintile indicated by the color legend among children born to parents whose income falls in the quintile indicated on the x-axis. We present these estimates along with those put forward by Chetty et al. (2014) for the United States (second panel) and Deutscher and Mazumder (2020) for Australia (third panel). While we rely on at most 11 income observations (7 on average) for parents and at most 7 income observations (5 on average) for children, Deutscher and Mazumder (2020) use 11 income observations for parents and 5 for children, and Chetty et al. (2014) use 5 income observations for parents and 2 for children.

We find that 9.7% of children born to parents in the bottom 20% reach the top 20% in

¹⁹Alesina et al. (2018) estimated father-son wage transition probabilities from the bottom quintile only, using the TSTSLs methodology and data from the *Formation et Qualification Professionnelle* survey for earlier cohorts (1963-1973).

their forties. This share is 7.5% in the United States and 12.3% in Australia. In comparison, 31.8% remain in the bottom 20% of the income distribution. Regarding children born to the top 20%, 38.4% remain at the top, while only 10.7% move down to the bottom of the income distribution, much less than in Australia (14%). As a reference point, in a society where an individual's income is completely independent of parent income, the probability of being in any quintile given a parent quintile would by definition be 20%. We analyze persistence at the top of the parent income distribution in more detail in Appendix Section C.5.

Note that among the corner elements of the transition matrix, the estimates of mobility (i.e., $P(\text{Child Top } 20\% \mid \text{Parent Bot. } 20\%)$ and $P(\text{Child Bot. } 20\% \mid \text{Parent Top } 20\%)$) are likely to be upper bounds, while estimates of persistence (i.e., $P(\text{Child Bot. } 20\% \mid \text{Parent Bot. } 20\%)$ and $P(\text{Child Top } 20\% \mid \text{Parent Top } 20\%)$) are likely to be lower bounds. This is because the potential measurement error in parent rank prediction induced by TSTSLS can only go in one direction for the bottom and top quintiles. Parents in the bottom 20% necessarily have a true rank in the bottom 20% or above, but not below, as ranks take positive values by definition. Reasonably assuming that the probability of reaching the top 20% is increasing in parent income rank, our estimate of $P(\text{Child Top } 20\% \mid \text{Parent Bot. } 20\%)$ is therefore likely to be an upper bound. In line with this intuition, the PSID validation exercise suggests that TSTSLS transition matrices overstate mobility relative to observed transition matrices (see Appendix Table B.4). The same reasoning can be applied to the other corner elements of the transition matrix.

In Table 1.2 we compare conditional probabilities of interest with those found for other developed countries. In France income persistence across generations is particularly strong, both at the top and at the bottom. While France does better than the United States when it comes to upward mobility from the bottom quintile (9.7% vs. 7.5%), a point we discuss in Section 4.4, it fares significantly worse than countries such as Canada (11.4%), Switzerland (11.9%) or Australia (12.3%). It also displays one of the strongest persistence at the bottom and at the top of the income distribution.

4.4. Discussion of Baseline Results

International Comparison. Our findings confirm the conventional wisdom that France exhibits strong income persistence across generations relative to many OECD countries (OECD, 2018). This is true not only with respect to the IGE, which has been the main focus for cross-country comparisons in the literature (e.g., see Corak (2016)), but also for the RRC, and in terms of transition matrices. This raises the question of the underlying mech-

Country	P(Child Top 20% Parent Bot. 20%) ↓	P(Child Bot. 20% Parent Bot. 20%)	P(Child Top 20% Parent Top 20%)	Source
United States	7.5%	33.7%	36.5%	Chetty et al. (2014, Table 2)
Italy ¹	8.6% ²	36.7%	27.8%	Acciari et al. (2022)
France	9.7%	31.8%	38.4%	
Denmark	10.7%	30.7%	34.8%	Eriksen (2018, Figure 3.3*)
Netherlands	11.3%	29.8%	33.1%	Carmichael et al. (2020, Table 1*)
Canada	11.4%	30.1%	32.3%	Corak (2020, Table 6)
Switzerland	11.9%	23.7%	30.3%	Chuard-Keller and Grassi (2021, Table 2)
Spain	12.3%	25.3%	33.3%	Soria Espín (2022, Table A.5)
Australia	12.3%	31%	30.7%	Deutscher and Mazumder (2020, Table 3)
Switzerland	12.8%	24.5%	28.8%	Kalambaden and Martinez (2021, Table 5)
Sweden ³	15.7%	26.3%	34.5%	Heidrich (2017, Figure 10, Appendix B)

Notes: See Table 1.1 for details about samples and income definitions used in each study.

¹ As the authors point out, this paper’s baseline estimates are likely to overestimate upward mobility and underestimate persistence at the bottom and at the top because of lifecycle bias, the omission of taxpayers and tax evasion. The reported P(Top 20% | Bottom 20%) here corresponds to the estimate accounting as best as possible for these three sources of bias. For the other two measures, we report the estimates correcting for missing tax returns and tax evasion obtained from the authors.

² Obtained by multiplying the “Q1Q5” estimate found in the last column of Table 14 by the ratio of the two rows in Table 11, i.e., $0.100 \times 0.099 / 0.115$.

³ Child incomes are measured relatively early in the lifecycle (32-34 years old), thus these estimates may suffer from lifecycle bias (i.e., overestimating upward mobility and underestimating persistence). By comparison, the father-son P(Child Top 20% — Parent Bot. 20%) estimate in Nybom and Stuhler (2017, Figure 1, Panel D) is essentially 10%, a much lower estimate of upward mobility.

* The authors very kindly shared more detailed estimates than reported in their papers.

Table 1.2: Transition Matrix in International Comparison

anisms. Indeed, one apparent puzzle is that various studies have found positive effects of government spending on intergenerational mobility (Mayer and Lopoo, 2008; Huang et al., 2021). Yet, despite significant government spending, France displays relatively little intergenerational mobility.

However, though the IGE and RRC estimates are fairly similar for France and the United States, the two countries differ in terms of the probability of reaching the top 20% conditional on having parents in the bottom 20%. Given the large dissimilarities in their higher education systems, part of the explanation could stem from differences in access to, and graduation from, higher education along the parent income distribution.

Access to and Graduation from Higher Education. Using the yearly census surveys available since 2004 in the EDP, we can observe children’s last obtained diploma when they are between 23 and 45.²⁰ Figure 1.5 compares higher education *graduation rates* in France with *enrollment rates* in the United States (defined by Chetty et al. (2020) as attending college at least at some point between ages 18-21) by parent income rank. To avoid capturing the direct effect of parent education (independent from parent income) on child

²⁰We observe this information for 86% of the sample. The share of missing values is pretty well uniformly distributed along the parent income rank distribution.

higher education graduation, we use parent income ranks obtained when excluding parent education from the set of first-stage predictors. This has virtually no effect on the result. Graduation rates in France are lower than enrollment rates in the United States, which is expected considering that a sizable share of students who enroll in higher education eventually drops out. While the relationship between parent income rank and enrollment is linear in the United States, obtaining a higher education degree appears to be a convex function of parent income rank in France. In particular, it is flatter at the bottom of the distribution.²¹ This convex relationship is all the more striking since children from low-income families are probably more likely to drop out from higher education, and therefore not earn a higher education degree.




Figure 1.5. Graduation From/Enrollment In Higher Education by Parent Income

Notes: This figure presents higher education graduation in France vs. enrollment rates in the United States (Chetty et al., 2020) by parent income rank. See Figure 1.3’s notes for details on data, sample and income definitions. In this figure parent income ranks are computed without parent education in the set of first-stage predictors to avoid capturing the effect of parent education independent from that of parent income.

This comparison does not allow us to assess directly whether higher education may explain the gap in upward mobility between France and the United States, since the rela-

²¹Appendix Figure E.6 documents the graduation rate for each cell of the quintile-by-quintile transition matrix. It shows that the convexity in the relationship between family background and graduation rate holds within child income quintile.

tionship between college completion and parent income rank for the latter is not available. Using a French survey of roughly 6,000 18-24 year olds, [Bonneau and Grobon \(2022\)](#) find that enrollment rates in higher education by parent income rank are very similar in France compared to the United States. Therefore, if higher education were to explain part of the upward mobility gap observed between the two countries, it must necessarily be through differences in dropout rates and/or heterogeneous returns to higher education along the parent income distribution.

5. Robustness of Baseline Results

In addition to the method validity exercise presented in Section 3.1, we assess the sensitivity of our baseline results to the TSTOLS method by (i) varying the set of instruments, and (ii) relaxing parametric assumptions. Moreover, as discussed in Section 2.2, two statistical biases may affect our baseline estimates: lifecycle and attenuation bias. The former relates to heterogeneous lifecycle earnings profiles among parents and children, while the latter refers to classical measurement error in parent income. We therefore assess how our estimates vary with the age at which child and parent incomes are measured, and with the number of synthetic parent income observations used. We discuss additional potential biases (i.e., data coverage, treatment of zero incomes, and top and bottom income trimming) in Appendix C.

5.1. Two-Sample Two-Stage Least Squares

First-Stage Predictors. We first estimate the IGE, RRC, and transition matrices using only education as the first-stage predictor. We then add successively to the set of first-stage predictors: parents' (i) 2-digit occupation, (ii) demographic characteristics, and (iii) municipality-level characteristics. Our baseline specification corresponds to the one including the full set of predictors. The results are shown in Appendix Figure C.2.

Overall, our estimates are largely insensitive to the set of first-stage regressors, except for the IGE which is significantly larger when using only education in the first-stage. For example, the RRC (IGE) when using only education is 0.284 (0.679) compared to 0.303 (0.527) in our baseline. The transition matrices are also mostly unchanged: when using only education the $P(\text{Top } 20\% - \text{Bot. } 20\%)$ is 10.8% compared to 9.7% in our baseline. These results are consistent with our validation exercise using the PSID where we find that the TSTOLS RRC estimate increases slightly once (3-digit) occupation is included as a predictor and the transition matrices are largely unaffected by the set of first-stage re-

gressors (see Appendix Table B.4).

Functional Form. We estimate the first-stage using the three following flexible methods: (i) generalized additive model (GAM), (ii) gradient boosted tree, and (iii) the ensemble method. The results are shown in Appendix Figure C.3. These more flexible models yield essentially identical estimates and they do not lead to gains in terms of (out-of-sample) mean squared error.

5.2. *Lifecycle and Attenuation Bias*

Child Lifecycle Bias. Figure 1.6 presents our estimates of intergenerational income mobility when varying the age at which child income is measured. In addition to household income from the tax returns data, we exploit the longer time series wage data provided by the All Employee Panel. Each point represents the estimate of the measure of intergenerational income mobility when measuring child income at a given age. For the transition matrix, we only present the analysis for the conditional probability of being in the top or bottom 20% for children born to parents in the top or bottom 20%. The broad pattern that emerges in Figure 1.6 is that the estimated persistence (mobility) increases (decreases) sharply when child incomes are measured early in the lifecycle and stabilizes roughly when child income is measured in their mid-thirties.²²

²²By construction, each age estimate is obtained from a different sample since we only measure child incomes in the tax returns data between 2010 and 2016, and in the All Employee Panel from 1967 to 2015 (though only for individuals born in even years before 2001). The observed slight decline in the IGE and RRC estimates when children are in their forties for household income appears to mostly reflect changes in the underlying cohort sample rather than a real decrease in the estimate (see Appendix Figure C.4 where we reproduce the All Employee Panel estimates keeping the sample of children constant).

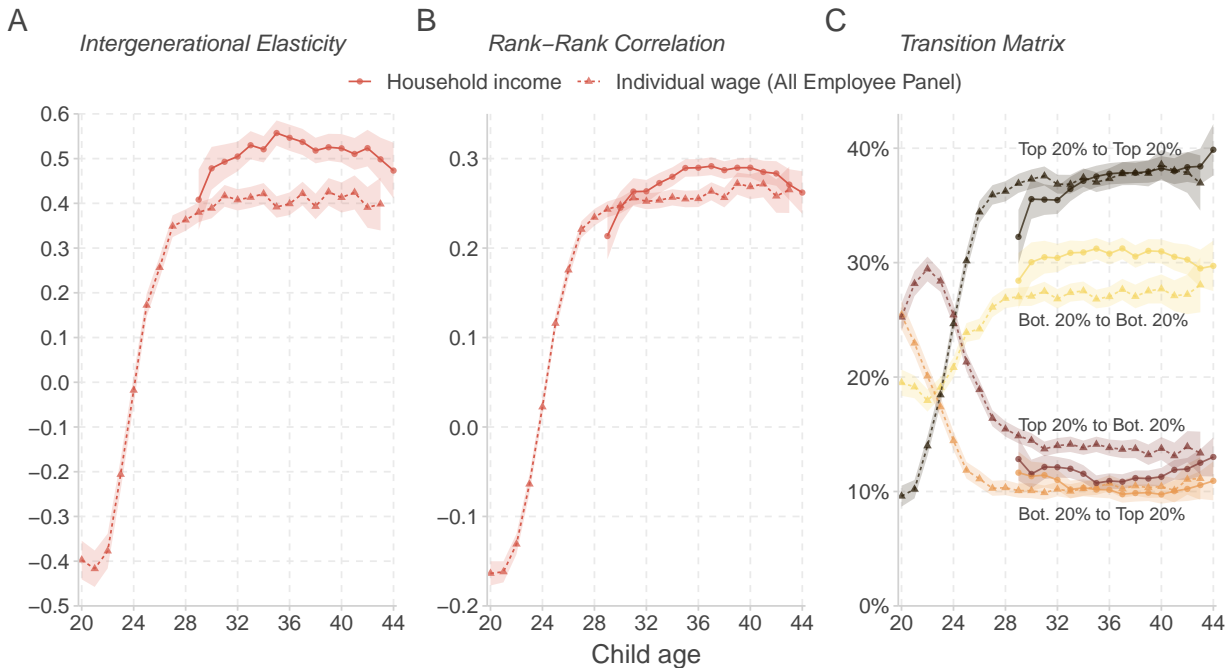


Figure 1.6. Child Lifecycle Bias

Notes: This figure assesses the robustness of our baseline intergenerational income mobility estimates to changes in the age at which child income is measured. Shaded areas represent the 95% bootstrapped confidence intervals. See Figure 1.3’s notes for details on data, sample and income definitions.

Parent Lifecycle Bias. We assess the sensitivity of our baseline estimates to varying the age at which parent income is measured. Since we predict parent income rather than observe it, we vary the age at which synthetic parent income is measured in the first-stage regression. Specifically, we run the first-stage regression (equation (1.3)) defining synthetic parent income at a given age between 25 and 60 years old. Figure 1.7 shows that the relationship between age at which parent income is measured and persistence is concave, strongly increasing between 25 and the late thirties and then stabilizing until the mid to late fifties. Relative to our baseline estimate, it does not appear that our choice of measuring synthetic parent income as the average between 35 and 45 years old is either too early or too late in the lifecycle.²³

Attenuation Bias. We evaluate the extent to which our baseline estimates are sensitive to the number of observations used to compute parent lifetime income. The main source

²³In Appendix Section C.3 we study how our measures of intergenerational persistence vary with the age at which child and synthetic parent income is measured jointly.

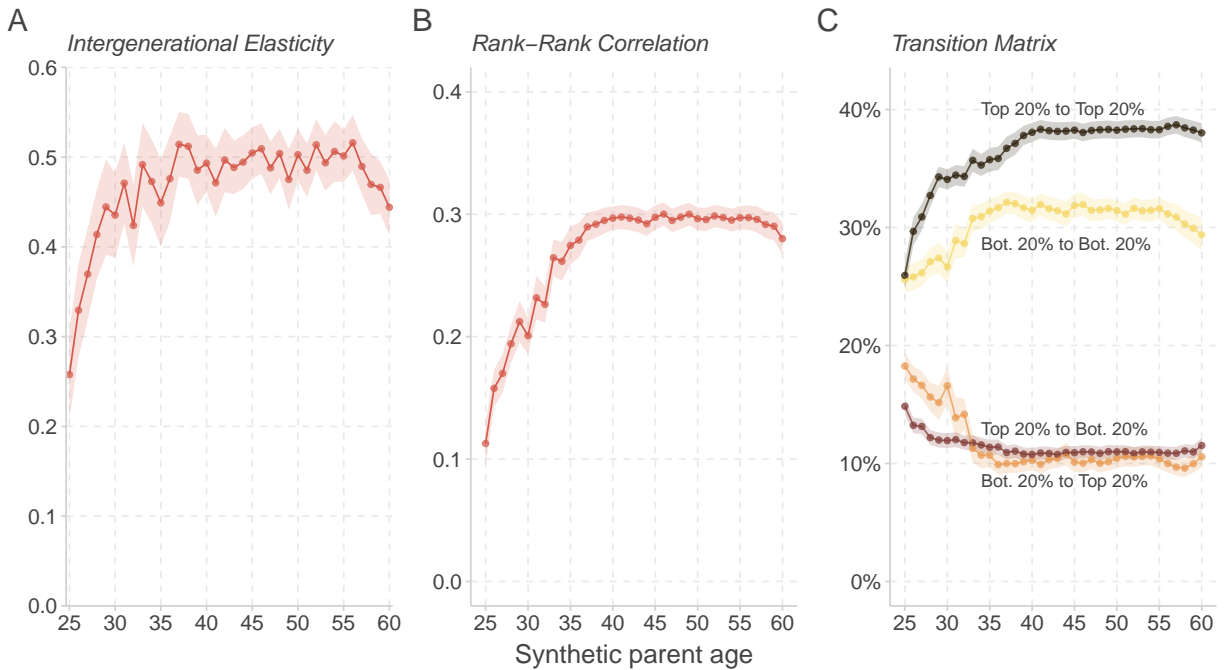


Figure 1.7. Parent Lifecycle Bias

Notes: This figure assesses the robustness of our baseline intergenerational income mobility estimates to changes in the age at which synthetic parent income is measured. Shaded areas represent the 95% bootstrapped confidence intervals. See Figure 1.3’s notes for details on data, sample and income definitions.

of attenuation bias comes from measurement error in parent income.²⁴ Appendix Figure C.6 plots estimates of our persistence measures varying the number of synthetic parent income observations used in the first-stage regression from 1 to 11 (see details in Appendix Section C.3). The rank-based measures, whether the RRC or the transition matrix cells, are remarkably unaltered by increasing the number of income observations over which synthetic parent income is averaged. However, the IGE increases gradually with the number of income observations, which largely rests on how mothers’ incomes are predicted. In the context of TSTSLS estimation, this appears to be a strength of rank-based measures since it suggests that in cases where parent income is not observed, predicting it using only one synthetic parent income observation is likely to provide sufficiently accurate estimates. This is indeed what we find in our validation exercise, where the TSTSLS RRC bias is largely unchanged when increasing the number of parent income observations.

²⁴We also check in Appendix Section C.3 the sensitivity of intergenerational mobility to the number of child income observations and confirm that it only plays a very minor role.

6. Geographic Analysis

6.1. Heterogeneity Across Departments

A first step in understanding the sources of intergenerational mobility in France is to investigate where persistence is highest and lowest. We study the geographic variations of intergenerational mobility at the department level. Departments divide metropolitan France into 96 territories.²⁵ Departments have the advantage of covering the whole of metropolitan France, and their borders have not changed over the study period. In addition, considering finer geographic units such as commuting zones would imply dropping a sizable amount of areas due to insufficient sample size.

Children are assigned to their department of residence in 1990, when they were between 9 and 18 years old. This is the best proxy we have for the department they grew up in. To ensure our estimates are sufficiently reliable, we focus on the 85 departments with at least 200 observations.²⁶ Individuals are still ranked within the national income distribution.

Hereinafter we use parent income predicted without municipality characteristics in the first stage. This is to make sure that they do not spuriously drive any spatial patterns.²⁷ Moreover, we find that spatial variations in intergenerational mobility are not driven by differences in prediction accuracy of the first-stage across departments. Indeed, as shown in Appendix Table F.8, the department-level mean-squared errors of the first-stage predictions are not significantly related with department-level intergenerational mobility measures.

The statistics we use at the subnational level are (i) the IGE, (ii) the RRC, and (iii) the expected income rank for individuals whose parents locate at the 25th percentile, which we refer to as *absolute upward mobility* (AUM) following Chetty et al. (2014). We favor absolute upward mobility over specific cells of the transition matrix because of the size of our department samples. Indeed, while absolute upward mobility is estimated using all the observations in a given department, any cell of the quintile transition matrix is by construction estimated using only a fifth of these observations. Denoting $p_{c,d}$ the percentile income rank of children observed in department d during childhood, and $p_{p,d}$ the

²⁵For practical reasons, we treat Corsica as a single department. Appendix Figure E.7 shows a map of French departments.

²⁶The number of observations per department is reported in Appendix Table F.9.

²⁷The removal of municipality characteristics from the first stage does not alter our national estimates (see Appendix Figure C.2) nor the first-stage R^2 . Moreover, the cross-department correlation with and without municipality characteristics is above 0.97 for all three intergenerational mobility measures (IGE, RRC, AUM).

percentile income rank of their parents, local RRCs are obtained from the following OLS regression:

$$p_{c,d} = \alpha_d + RRC_d \times p_{p,d} + \varepsilon_d \quad (1.4)$$

The expected income rank for individuals whose parents locate at the 25th percentile then writes:

$$AUM := \mathbb{E}[p_{c,d} \mid p_{p,d} = 25] = \hat{\alpha}_d + R\hat{R}C_d \times 25 \quad (1.5)$$

Appendix Figure E.8 graphically illustrates how this intergenerational mobility measure is computed for the Nord department, the most populated one in 1990. The conditional expectation functions for the most populated departments are available in Appendix Figures E.9 and E.10. Even at the department level, it appears that the rank-rank relationship is well approximated by a linear function.

Geographic Variations. Figure 1.8 depicts department-level intergenerational mobility as captured by the three estimators mentioned above. It reveals substantial variations, though not necessarily statistically significant likely due to a lack of statistical power.²⁸ The distribution of department-level RRCs ranges from 0.17 to 0.40 and is tighter than that of IGEs, which ranges from 0.27 to 0.88. Both vary across departments just as much as they vary across countries. The range of our estimates of absolute upward mobility, from rank 37 to rank 54, is almost identical to that observed in Italy using a comparable geographic unit (from 35 to 57 (Acciari et al., 2022)).

Intergenerational persistence is particularly high in the North and in the South of France, and relatively low in the West. For instance, the IGEs range from 0.30 to 0.45 in departments in Brittany (West), from 0.42 to 0.70 in departments in Hauts-de-France (North), and from 0.63 to 0.77 in the former region of Languedoc-Roussillon (South). This pattern is observed not only in terms of relative mobility (IGE and RRC), but also in terms of absolute upward mobility. Indeed, while children with modest socio-economic backgrounds have relatively high expected income ranks in Brittany (AUM \in (43.3; 44.7)), they tend to remain lower in the income distribution in Hauts-de-France (AUM \in (36.8; 44.1)) and Languedoc-Roussillon (AUM \in (36.9; 39.3)).

However, a high relative mobility is not systematically associated with a high absolute upward mobility. For instance, such a discrepancy is observed for the municipality-

²⁸Department-level estimates are reported in Appendix Table F.9. Department-level IGE, RRC and AUM are represented graphically with their confidence intervals in Appendix Figures E.11 to E.13.

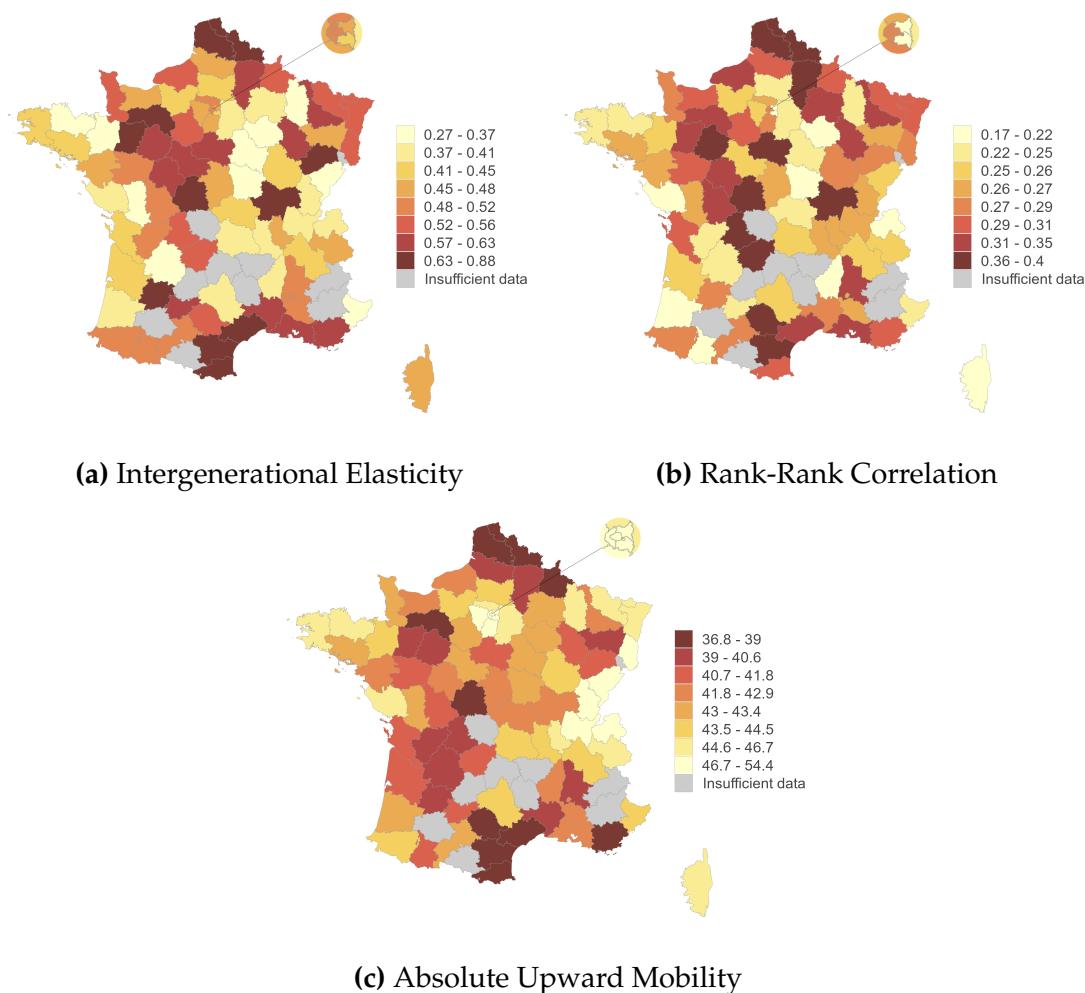


Figure 1.8. Spatial Variations in Intergenerational Mobility

Notes: This figure presents department-level estimates of our intergenerational mobility measures. To compute local estimates, individuals are assigned to their department of residence in 1990, when they were between 9 and 18 years old. Departments with less than 200 observations are considered as having insufficient data. See Figure 1.3’s notes for details on data, sample and income definitions.

department of Paris, the third highest department in terms of AUM, but where intergenerational mobility levels in terms of IGE and RRC are close to the department-level average. The conditional expectation functions in Appendix Figure E.10 provide an explanation to this idiosyncrasy. They reveal that the Parisian CEF is both shifted upwards relative to other large departments, and flatter at the lower end of the parent income distribution. The combination of these two features results in relatively good prospects for children whose parents locate at the 25th percentile without implying particularly high relative mobility. The cross-department correlation between the IGE and RRC is 0.65, and is -0.55 with AUM (see Appendix Table F.10), which highlights the importance of using

a variety of intergenerational mobility measures to characterize a country's income persistence across generations (Deutscher and Mazumder, forthcoming).

Correlation with Local Characteristics. To pin down potential sources of the spatial variations in intergenerational mobility, we explore the department characteristics that it might correlate with. We consider 14 variables, measured as close to 1990 as possible, classified into 5 groups: demographic, economic, inequality, education, and social capital variables. There are three main takeaways from this correlational analysis (additional details can be found in Appendix D).

First, the IGE appears to be only significantly related to the unemployment rate. This correlation is indeed striking visually when comparing the department-level unemployment rate in 1990, displayed in Appendix Figure E.14, with Figure 1.8a. Second, absolute upward mobility tends to exhibit much stronger relationships with department characteristics in general, than either the IGE or the RRC. This suggests that factors that affect absolute mobility might differ from those that affect relative mobility. A lasso analysis, detailed in Appendix D.3, yields similar insights.

Third, we find no evidence of a within France “Great Gatsby Curve”, which refers to the positive correlation between intergenerational income persistence (defined by the IGE) and income inequality (defined by the Gini index) found across countries (Corak, 2013). The Gini index is significantly positively related to absolute upward mobility, the opposite sign one might expect if inequality is detrimental to intergenerational mobility. This contrasts with findings from Italy (Acciari et al., 2022) and North America (Chetty et al. (2014) for the United States and Corak (2020) for Canada).

6.2. *Geographic Mobility*

Few studies have explored the relationship between geographic mobility and intergenerational mobility.²⁹ We consider individuals as geographically mobile if their adulthood department of residence is different from their childhood department of residence. The childhood department of residence is observed in the 1990 census, when individuals were aged from 9 to 18 years old. The adulthood department of residence is the one indicated on individuals' tax return. If the individual has lived in several departments over 2010-2016, we consider the most common department of residence. In case of ties, we consider

²⁹Soria Espín (2022) analyzes this relationship in Spain, but other existing studies rather exploit geographic mobility to estimate the causal impact of location on upward mobility (Chetty and Hendren, 2018; Laliberté, 2021).




Figure 1.9. Intergenerational Mobility and Department Characteristics - Separate Estimation

Notes: This figure presents the regression coefficient between department-level intergenerational mobility and department characteristics. Each coefficient is obtained from a separate regression. Both the department intergenerational mobility estimates and the characteristics are standardized, implying that the coefficients can be interpreted as correlations. Horizontal lines represent the 95% confidence intervals. See Figures 1.3 and 1.8's notes for details on data, sample and income definitions, and Appendix Table D.1 for definitions and sources of the department characteristics.

the most recent of the most common departments. According to this definition, 40.8% of individuals are geographically mobile. This share is relatively homogeneous across males (40.2%) and females (41.3%). The percentage of movers by parent household wage rank is presented in Appendix Figure E.15.

Intergenerational Mobility Gains from Geographic Mobility. Figure 1.10 shows the conditional expectation of child household income rank with respect to parent household wage rank for movers and stayers. The CEF is slightly flatter for movers than for stayers, and importantly, movers have systematically higher expected income ranks than stayers throughout the parent household wage rank distribution. The difference between the two CEFs is slightly decreasing in parent income and is particularly pronounced at the bottom of the distribution. This difference is the result of the combination of individuals self-selecting into migration and the causal effect of moving.



Figure 1.10. Intergenerational Mobility and Geographic Mobility

Notes: This figure represents the conditional expectation of child household income rank with respect to parent household wage rank separately for individuals whose adulthood department of residence is different or not from their childhood department of residence. Percentile ranks are computed according to the national income distribution, which implies that the share of movers and stayers is not constant throughout the parent income distribution. The childhood department of residence is observed in the 1990 census, when individuals were aged from 9 to 18 years old. The adulthood department of residence is the one indicated on individuals' tax return. If the individual has lived in several departments over 2010-2016, we consider the most represented department of residence. In case of ties, we consider the most recent of the most represented departments. See Figure 1.3's notes for details on data, sample and income definitions.

To characterize the relationship between intergenerational and geographic mobility, we estimate the following regression model:

$$p_{c,i} = \alpha + \beta p_{p,i} + \gamma \text{Mover}_i + \delta p_{p,i} \times \text{Mover}_i + X_i' \lambda + \varepsilon_i, \quad (1.6)$$

where $p_{c,i}$ is the household income rank of individual i , $p_{p,i}$ is individual i 's parents' household wage rank, Mover_i is a binary variable taking the value 1 if individual i lives in a different department from the one they grew up in and 0 otherwise, and X_i is a set of control variables. Table 1.3 reports the corresponding regression results.

Column (1) shows the estimates from equation (1.6). Living in a different department from one's childhood department is associated, on average, with a $\mathbb{E}[\hat{\gamma} + \hat{\delta} p_{p,i}] = 5.89$ percentile rank increase in the national household income distribution. The point estimate of

	<i>Dependent variable: Child household income rank</i>				
	(1)	(2)	(3)	(4)	(5)
Parent income rank ($\hat{\beta}$)	0.278*** (0.005)	0.278*** (0.005)	0.268*** (0.005)	0.163*** (0.008)	0.138*** (0.017)
Mover ($\hat{\gamma}$)	5.836*** (0.472)	5.858*** (0.472)	5.539*** (0.475)	5.716*** (0.472)	5.681*** (0.475)
Parent income rank \times Mover ($\hat{\delta}$)	0.001 (0.008)	0.0003 (0.008)	0.001 (0.008)	-0.012 (0.008)	-0.013 (0.008)
Constant	34.087*** (0.258)	33.780*** (0.274)	38.123*** (1.228)	29.195*** (1.659)	30.509*** (1.782)
Birth cohort	✓	✓	✓	✓	✓
Gender		✓	✓	✓	✓
Department FE			✓	✓	✓
Parents' education				✓	✓
Parents' 2-digit occupation					✓
$\mathbb{E}[\hat{\gamma} + \hat{\delta}p_p] = \hat{\gamma} + \hat{\delta} \times 50.5$	5.89	5.87	5.59	5.11	5.02
$\mathbb{E}[\hat{\gamma} + \hat{\delta}p_p p_p = 25]$	5.86	5.86	5.56	5.42	5.36
$\mathbb{E}[\hat{\gamma} + \hat{\delta}p_p p_p = 75]$	5.91	5.91	5.61	4.82	4.71
Observations	64,571	64,571	64,571	64,571	64,571
Adjusted R ²	0.098	0.098	0.106	0.119	0.125

Notes: This table provides the estimates from regression child household income rank on their parents' income rank, a dummy variable indicating whether the individual is a mover, and the interaction between these two variables. Columns (2) to (5) progressively include control variables. See Figure 1.10 for details on variable and sample definitions. Bootstrapped standard errors in parentheses. *p<0.1; **p<0.05; ***p<0.01.

Table 1.3: Intergenerational & Geographic Mobility

the rank-rank slope is slightly lower for movers when controlling for parents characteristics (coefficient $\hat{\delta}$ col. (4)-(5)), but not statistically significantly so. In the last specification, the difference in expected income rank between movers and stayers is decreasing in parent income (5.36 at the 25th percentile and 4.71 at the 75th percentile).

The Role of Mobility Toward Richer Departments at the Aggregate Level. There are several potential reasons for the better intergenerational mobility outcomes movers tend to experience. One explanation may be that movers simply migrate to departments where wages are higher. To investigate this channel, we compute two statistics: (i) the mean parent household wage rank in the origin department, and (ii) the mean child household income rank in the destination department. Figure 1.11 displays the average of these two

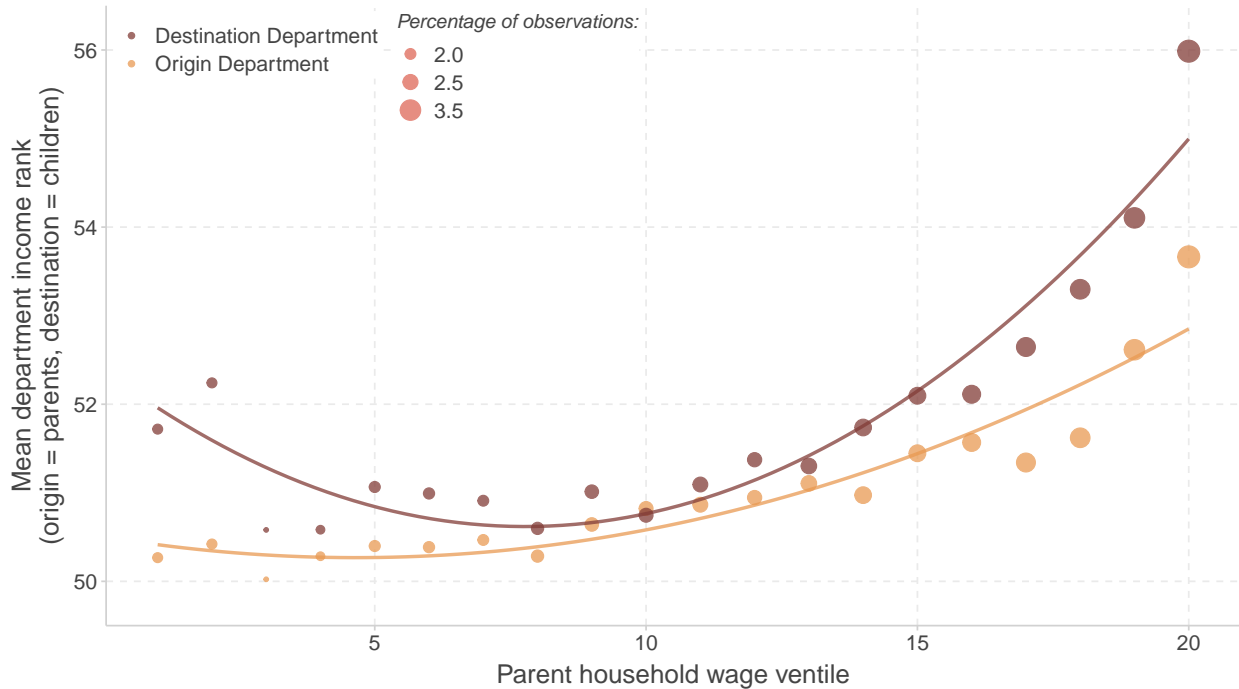


Figure 1.11. Mean Income Rank of Origin and Destination Departments of Movers

Notes: This figure represents the conditional expectation of income rank with respect to parent household wage rank for movers, separately by origin and destination departments. Origin department mean income rank is computed as the average income rank of residents in the parent sample, while destination mean income rank is computed as the average income rank of residents in the child sample. See Figures 1.3 and 1.10’s notes for details on data, sample and income definitions.

statistics for movers for each ventile of the parent household wage rank distribution.

There are three takeaways from this figure. First, the difference in average income rank in the destination and origin departments is highest at the top and bottom of the parent income distribution. Second, these differences are relatively small, reaching at most 2 percentile ranks for the top ventile. Third, the origin and destination departments of movers from the middle of the parent income distribution have very similar average income ranks. Put in parallel with the slight monotonic decrease in the gains from geographic mobility along the parent income rank distribution, it seems that these gains are not only due to individuals moving to higher-income departments.

Another way to test this hypothesis consists in comparing the conditional expectation functions of movers and stayers ranked either at the *national* and *department* level. Indeed, ranking individuals at the national level allows individuals born to parents who earn the median income of their department to be upward mobile by earning the median income of a higher-income department in adulthood. This channel can be removed

by ranking individuals and their parents within departments. When doing so, movers can only be more intergenerationally mobile than stayers if they reach income ranks in their adulthood department that are further away from the rank of their parents in their childhood department. Finding no expected gains associated with geographic mobility when ranking individuals according to their department income distribution would suggest that the expected increase in income rank associated with mobility is fully driven by movers ending up in higher-income departments, but reaching on expectation a local income rank in their destination department that is not further away from that of their parents, relative to stayers.

The regression results of equation (1.6) using percentile ranks computed at the department level rather than at the national level are reported in Appendix Table F.11 (Appendix Figure E.16 shows the corresponding conditional expectation functions). When considering ranks in the department distribution, the gap between the conditional expectation functions of movers and stayers shrinks but does not vanish completely. While the expected national-rank increase associated with mobility amounts to 5.89, it drops to 3.87 when considering local ranks. This suggests that the intergenerational mobility gains associated with geographic mobility are partly attributable to movers locating in higher-income departments in adulthood relative to stayers, but also to movers reaching local ranks in their adulthood department that are further away from the rank of their parents in the childhood department.

The Role of Mobility Toward Richer Departments at the Individual Level. While geographic mobility patterns between low- and high-income departments only partially explain the gap between movers and stayers at the aggregate level, characteristics of the destination department may be decisive at the individual level. To investigate this hypothesis we classify destination departments into three groups according to the average income rank of their residents from the child sample: (i) *low-income*, destination departments with an average income rank below 50 (70 departments - 49% of movers), (ii) *medium-income*, those with an average income rank between 50 and 60 (20 departments and overseas departments - 33% of movers), and (iii) *high-income*, those with an average income rank above 60 (5 departments and foreign countries - 18% of movers). This high-income group of departments greatly overlaps with the Parisian region as it comprises Essonne, Hauts-de-Seine, Paris, and Yvelines.

Figure 1.12 shows the conditional expectation of child income rank with respect to parent income ventile for the three destination department categories and for stayers. Results

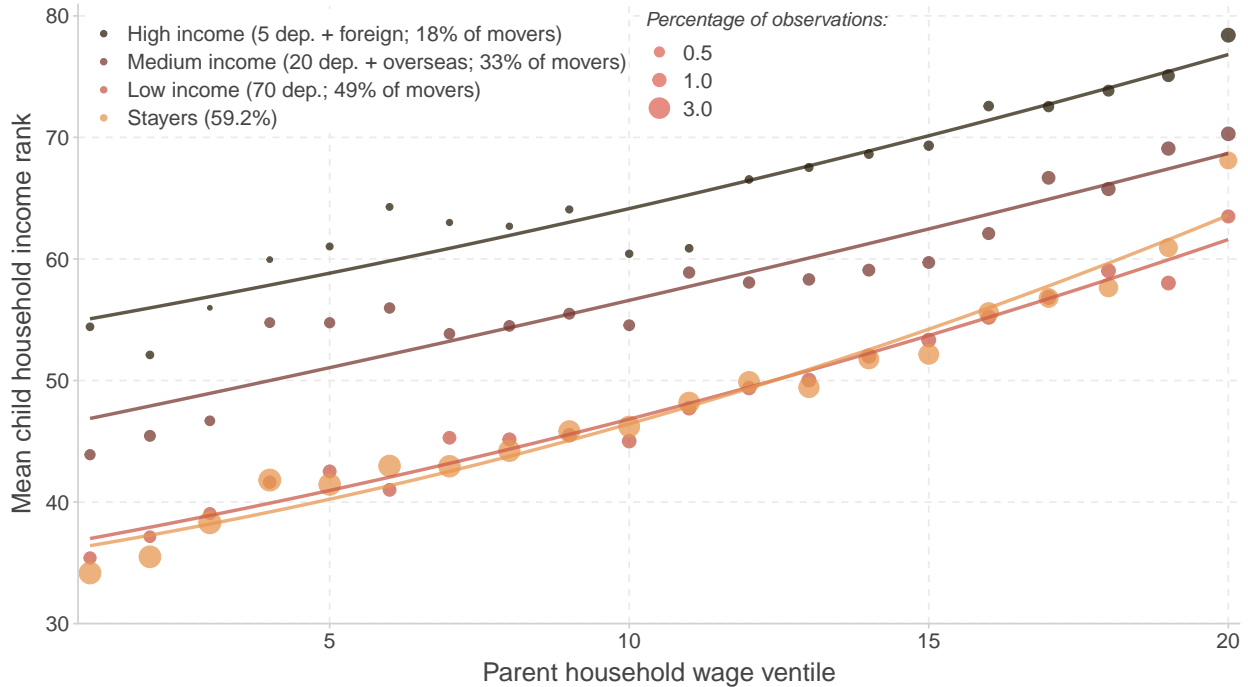


Figure 1.12. Mean Child Income Rank by Destination Department Mean Income

Notes: This figure represents the conditional expectation of child household income rank with respect to parent household wage rank for stayers and for movers to departments of different mean income categories. Solid lines represent second-order polynomial fits. Low income destination departments are destination departments with an average income rank below 50, medium income are those with an average income rank between 50 and 60, and high income are those with an average income rank above 60. See Figures 1.3 and 1.10’s notes for details on data, sample and income definitions.

of the corresponding regression are reported in Appendix Table F.12. Except for the top ventiles, the CEFs of movers by destination department category are virtually parallel. Movers thus experience similar levels of relative mobility regardless of the income category of their destination department. However, movers’ absolute upward mobility increases with the average income of the destination department, such that the expected income rank of a mover from the bottom of the parent income distribution to a high-income department is around the same as the expected income rank of a stayer from the 75th percentile of the parental income distribution. Still, such transitions are the exception: most movers to high-income departments come from high-income families, while low-income movers go predominantly to low- or medium-income departments. Another noteworthy finding is that expected income ranks are essentially the same for movers to low-income departments as for stayers, highlighting the potential role of the destination department’s characteristics in generating upward intergenerational mobility for movers. All these findings combine self-selection and causal effects, and we leave the

disentangling of these two channels for future research.

7. Conclusion

France is an interesting case study for intergenerational income mobility considering its relatively modest income inequality and the specificity of its higher education system. Yet, it has been the focus of few studies due to important data limitations. We use administrative data to provide an overview of intergenerational income mobility in France for individuals born in 1972-1981. Relative to existing studies, the richness of these data enables us to apply two-sample two-stage least squares (TSTOLS) using a much larger set of individual characteristics, and to extensively assess the robustness of the resulting estimates. Using the Panel Study of Income Dynamics (PSID) we find that the TSTOLS methodology slightly underestimates rank-based measures of intergenerational persistence relative to what would be obtained if parent income was observed.

Moreover, we provide the first estimates of the rank-rank correlation and transition matrix for France, and conduct a comparative analysis with other countries for which such statistics are available. Our results reveal that France exhibits a relatively strong intergenerational income persistence at the national level. It ranks among the highest in OECD countries, with Italy and the United States, and far from Australia, Canada, and Scandinavian countries.

This high intergenerational income persistence at the national level hides substantial geographic heterogeneity across departments. We observe about as much variation across French departments as across countries. Intergenerational persistence appears to be particularly high in the North and South, and relatively low in the Western part of the country. Yet, only *absolute* mobility, as opposed to *relative* mobility, significantly correlates with local characteristics.

We also provide novel descriptive evidence on a new mechanism that could explain some features of intergenerational mobility: geographic mobility. We find that the difference in expected income ranks between geographically mobile individuals and stayers is large and slightly decreasing in parent income. This difference appears not to be solely due to individuals moving to higher income departments but to be also the result of individuals moving up the local income rank ladder. Destination departments are on average characterized by higher income levels than origin departments only at the tails of the parent income distribution. However, regardless of parent income rank, conditional on moving the absolute upward mobility gains associated with moving to a higher-income

department appear to be large and increasing with average income in the destination department. Even though not causal, we believe that these descriptive findings constitute promising avenues for future research to better understand intergenerational income mobility.

References

- Abbas, Hicham and Michaël Sicsic**, “Qui Gravit l’échelle Des Revenus Par Rapport à Ses Parents ?,” Documents de Travail N° 2022 - F04, INSEE, Paris 2022.
- Acciari, Paolo, Alberto Polo, and Giovanni L. Violante**, ““And Yet, It Moves”: Intergenerational Mobility in Italy,” *American Economic Journal: Applied Economics*, July 2022, 14 (3), 118–163.
- Alesina, Alberto, Stefanie Stantcheva, and Edoardo Teso**, “Intergenerational Mobility and Preferences for Redistribution,” *American Economic Review*, February 2018, 108 (2), 521–554.
- Barbieri, Teresa, Francesco Bloise, and Michele Raitano**, “Intergenerational Earnings Inequality: New Evidence From Italy,” *Review of Income and Wealth*, 2020, 66 (2), 418–443.
- Bellégo, Christophe, David Benatia, and Louis-Daniel Pape**, “Dealing with Logs and Zeros in Regression Models,” Technical Report, SSRN Electronic Journal 2021.
- Björklund, Anders and Markus Jäntti**, “Intergenerational Income Mobility in Sweden Compared to the United States,” *American Economic Review*, 1997, 87 (5), 1009–1018.
- , **Jesper Roine, and Daniel Waldenström**, “Intergenerational Top Income Mobility in Sweden: Capitalist Dynasties in the Land of Equal Opportunity?,” *Journal of Public Economics*, June 2012, 96 (5), 474–484.
- Black, Sandra E. and Paul J. Devereux**, “Recent Developments in Intergenerational Mobility,” in “Handbook of Labor Economics,” Vol. 4, Elsevier, 2011, pp. 1487–1541.
- Bloise, Francesco, Paolo Brunori, and Patrizio Piraino**, “Estimating Intergenerational Income Mobility on Sub-Optimal Data: A Machine Learning Approach,” *The Journal of Economic Inequality*, December 2021, 19 (4), 643–665.
- Bonneau, Cécile and Sébastien Grobon**, “Unequal Access to Higher Education Based on Parental Income: Evidence from France,” Technical Report 2022/01, World Inequality Lab, Paris 2022.
- Bratberg, Espen, Jonathan Davis, Bhashkar Mazumder, Martin Nybom, Daniel D. Schnitzlein, and Kjell Vaage**, “A Comparison of Intergenerational Mobility Curves in Germany, Norway, Sweden, and the US,” *The Scandinavian Journal of Economics*, 2017, 119 (1), 72–101.
- Carmichael, Fiona, Christian K. Darko, Marco G. Ercolani, Ceren Ozgen, and W. Stanley Siebert**, “Evidence on Intergenerational Income Transmission Using Complete Dutch Population Data,” *Economics Letters*, 2020, 189, 108996.

- Chetty, Raj and Nathaniel Hendren**, “The Impacts of Neighborhoods on Intergenerational Mobility II: County-Level Estimates,” *The Quarterly Journal of Economics*, August 2018, 133 (3), 1163–1228.
- , **John N Friedman, Emmanuel Saez, Nicholas Turner, and Danny Yagan**, “Income Segregation and Intergenerational Mobility Across Colleges in the United States,” *The Quarterly Journal of Economics*, August 2020, 135 (3), 1567–1633.
- , **Nathaniel Hendren, Patrick Kline, and Emmanuel Saez**, “Where Is the Land of Opportunity? The Geography of Intergenerational Mobility in the United States,” *The Quarterly Journal of Economics*, 2014, 129 (4), 1553–1623.
- Chuard-Keller, Patrick and Veronica Grassi**, “Switzer-Land of Opportunity: Intergenerational Income Mobility in the Land of Vocational Education,” Technical Report, SSRN Electronic Journal 2021.
- Corak, Miles**, “Income Inequality, Equality of Opportunity, and Intergenerational Mobility,” *Journal of Economic Perspectives*, 2013, 27 (3), 79–102.
- , “Inequality from Generation to Generation: The United States in Comparison,” Discussion Paper 9929, IZA, Bonn, Germany 2016.
- , “The Canadian Geography of Intergenerational Income Mobility,” *The Economic Journal*, 2020, 130 (631), 2134–2174.
- Cortes-Orihuela, Javier, Juan Díaz, Pablo Gutierrez Cubillos, Pablo Troncoso, and Gabriel Villarroel**, “How Much Should We Trust TSTOLS Intergenerational Mobility Estimates?: Evidence From A Developing Country,” *SSRN Electronic Journal*, 2022.
- Couch, Kenneth A. and Dean R. Lillard**, “Sample Selection Rules and the Intergenerational Correlation of Earnings,” *Labour Economics*, 1998, 5 (3), 313–329.
- Dahl, Molly and Thomas DeLeire**, “The Association between Children’s Earnings and Fathers’ Lifetime Earnings: Estimates Using Administrative Data,” Technical Report No. 1342-08, Institute for Research on Poverty 2008.
- Deutscher, Nathan and Bhashkar Mazumder**, “Intergenerational Mobility across Australia and the Stability of Regional Estimates,” *Labour Economics*, 2020, 66, 101861.
- and – , “Measuring Intergenerational Income Mobility: A Synthesis of Approaches,” *Journal of Economic Literature*, forthcoming.
- Eriksen, Jesper**, “Finding the Land of Opportunity: Intergenerational Mobility in Denmark,” Technical Report, Aalborg University 2018.
- Espín, Javier Soria**, “Intergenerational Mobility, Gender Differences and the Role of Out-Migration: New Evidence from Spain,” 2022, p. 98.
- Grawe, Nathan D.**, “Lifecycle Bias in Estimates of Intergenerational Earnings Persistence,” *Labour Economics*, 2006, 13 (5), 551–570.

- Haider, Steven and Gary Solon**, "Life-Cycle Variation in the Association between Current and Lifetime Earnings," *American Economic Review*, 2006, 96 (4), 1308–1320.
- Heidrich, Stefanie**, "Intergenerational Mobility in Sweden: A Regional Perspective," *Journal of Population Economics*, 2017, 30 (4), 1241–1280.
- Helsø, Anne-Line**, "Intergenerational Income Mobility in Denmark and the United States," *The Scandinavian Journal of Economics*, 2021, 123 (2), 508–531.
- Huang, Xiao, Shoujun Huang, and Ailun Shui**, "Government Spending and Intergenerational Income Mobility: Evidence from China," *Journal of Economic Behavior & Organization*, November 2021, 191, 387–414.
- Inoue, Atsushi and Gary Solon**, "Two-Sample Instrumental Variables Estimators," *The Review of Economics and Statistics*, 2010, 92 (3), 557–561.
- Jacome, Elisa, Ilyana Kuziemko, and Suresh Naidu**, "Mobility for All: Representative Intergenerational Mobility Estimates over the 20th Century," 2023.
- Jerrim, John, Álvaro Choi, and Rosa Simancas**, "Two-Sample Two-Stage Least Squares (TSTLS) Estimates of Earnings Mobility: How Consistent Are They?," *Survey Research Methods*, 2016, 10 (2), 85–102.
- Kalambaden, Preetha and Isabel Z Martinez**, "Intergenerational Mobility Along Multiple Dimensions Evidence from Switzerland," 2021.
- Laliberté, Jean-William**, "Long-Term Contextual Effects in Education: Schools and Neighborhoods," *American Economic Journal: Economic Policy*, May 2021, 13 (2), 336–377.
- Landersø, Rasmus and James J. Heckman**, "The Scandinavian Fantasy: The Sources of Intergenerational Mobility in Denmark and the US," *The Scandinavian Journal of Economics*, 2017, 119 (1), 178–230.
- Lee, Chul-In and Gary Solon**, "Trends in Intergenerational Income Mobility," *Review of Economics and Statistics*, 2009, 91 (4), 766–772.
- Lefranc, Arnaud**, "Intergenerational Earnings Persistence and Economic Inequality in the Long Run: Evidence from French Cohorts, 1931-75," *Economica*, 2018, 85 (340), 808–845.
- **and Alain Trannoy**, "Intergenerational Earnings Mobility in France: Is France More Mobile than the US?," *Annales d'Économie et de Statistique*, 2005, (78), 57–77.
- **, Nicolas Pistoiesi, and Alain Trannoy**, "Equality of Opportunity and Luck: Definitions and Testable Conditions, with an Application to Income in France," *Journal of Public Economics*, December 2009, 93 (11-12), 1189–1207.
- Mayer, Susan E. and Leonard M. Lopoo**, "Government Spending and Intergenerational Mobility," *Journal of Public Economics*, February 2008, 92 (1-2), 139–158.

Mazumder, Bhashkar, “Fortunate Sons: New Estimates of Intergenerational Mobility in the United States Using Social Security Earnings Data,” *Review of Economics and Statistics*, 2005, 87 (2), 235–255.

– , “Estimating the Intergenerational Elasticity and Rank Association in the United States: Overcoming the Current Limitations of Tax Data,” in Lorenzo Cappellari, Solomon W. Polachek, and Konstantinos Tatsiramos, eds., *Research in Labor Economics*, Vol. 43, Emerald Group Publishing Limited, 2016, pp. 83–129.

Nybohm, Martin and Jan Stuhler, “Heterogeneous Income Profiles and Lifecycle Bias in Intergenerational Mobility Estimation,” *Journal of Human Resources*, January 2016, 51 (1), 239–268.

– **and** – , “Biases in Standard Measures of Intergenerational Income Dependence,” *Journal of Human Resources*, 2017, 52 (3), 800–825.

OECD, *A Broken Social Elevator? How to Promote Social Mobility*, OECD, June 2018.

O’Neill, Donal, Olive Sweetman, and Dirk Van de gaer, “The Effects of Measurement Error and Omitted Variables When Using Transition Matrices to Measure Intergenerational Mobility,” *Journal of Economic Inequality*, 2007, 5 (2), 159–178.

Robert-Bobée, Isabelle and Natacha Gualbert, “L’échantillon Démographique Permanent : En 50 Ans, l’EDP a Bien Grandi !,” in “*Courrier Des Statistiques N6*” number 6, Paris: INSEE, 2021.

Solon, Gary, “Intergenerational Income Mobility in the United States,” *American Economic Review*, 1992, 82 (3), 393–408.

Zimmerman, David J, “Regression Toward Mediocrity in Economic Stature,” *American Economic Review*, 1992, 82 (3), 409–429.

A. Data - Details

The Permanent Demographic Sample (EDP) is a panel of individuals which the French statistical office, INSEE, started in 1968.³⁰ It combines several administrative data sources on individuals born on the first four days of October.³¹ Individuals born on one of these days are called EDP individuals. The EDP gathers data from 5 administrative sources: (i) civil registers since 1968; (ii) population censuses since 1968 (exhaustive in 1968, 1975, 1982, 1990 and 1999, and yearly rotating 20% random samples since 2004); (iii) the electoral register since 1990; (iv) the *All Employee Panel* since 1967; and (v) tax returns since fiscal year 2011.

Each time an individual born on the first four days of October appears in one of these administrative datasets, the information contained in it is added to their individual identifier in the EDP. Therefore all these datasets can be matched together using a common individual identifier. For our analysis we use data from civil registers, the 1990 census, the All Employee Panel and tax returns. We describe each data source in detail below.

Civil Registers. They contain information from birth certificates of EDP individuals and their children, as well as death and marriage certificates of EDP individuals, since 1968. We use birth certificates of EDP individuals and their children which include the child's gender, date and place of birth, and information on each parent including date and place of birth, nationality and occupation. There are no data breaks or missing certificates for the years under study (1972-1981).

1990 Census. It contains socio-demographic information about EDP individuals, as well as, though to a lesser extent, about members of their household. These include the individual's date and place of birth, nationality, education, occupation, marital status, household structure, dwelling characteristics, building when relevant, and municipality.

All Employee Panel. It combines two sources of data: the annual declarations of social data (*déclarations annuelles des données sociales* - DADS) and data on central government employees (*fichiers de paie des agents de l'état* - FPE). All businesses are obliged to annually communicate the declarations of social data about their employees to a network of private organizations (*Unions de recouvrement des cotisations de sécurité sociale et d'allocations familiales* - URSSAF) coordinated by a government agency (*Agence centrale des organismes de sécurité sociale* - ACOSS). The All Employee Panel data are reported at the worker-year level, aggregated by INSEE from data at the worker-firm-year level. As such, annual pretax wage and annual hours worked correspond to the sum over all the individual's salaried activities. The job characteristics correspond to the year's "main" job, that is the job for which the pay period was the longest and, in case of a tie, the job with the highest wage.

Between 1967 and 2001, data is only available for individuals born on an even year. The scope of workers covered by the All Employee Panel has varied over time. Since 1967 in metropolitan France, all private sector employees, except those in the agricultural sectors, and including employees of public enterprises, are covered. The hospital public service is integrated in 1984, the

³⁰The EDP user guide (in French) can be found [here](#).

³¹The EDP selection criterion has progressively widened to include individuals born on the first days of January, April, and July. See [Robert-Bobée and Gualbert \(2021\)](#) for a detailed description of the dataset.

state civil service and local authorities in 1988. France Télécom and La Poste employees appear only in 1988 as well. See Appendix C.1 for a robustness check to this public sector coverage evolution. The agricultural sector and overseas territories are included in 2002, and employees of private employers in 2009. Unemployment insurance is included from 2008 onwards. Lastly, because of increased workload due to the population censuses of 1982 and 1990, the All Employee Panel data were not compiled by INSEE in 1981, 1983 and 1990.

Tax Returns. They are compiled using housing and income tax forms filed for incomes earned from 2010 to 2016. In particular, household-level tax returns information is constructed based on dwellings where an EDP individual is known either from the income tax return or from the principal housing tax (*taxe d'habitation principale*). The location of the individual is that declared on January 1st of the fiscal declaration year. Income variables are available at the household-level as well as at the individual level. Since the information is gathered based on living in the same dwelling, household income is computed not only for couples who file their taxes jointly, but also for couples who live together, an increasingly common arrangement. This departs from existing studies based on tax returns data which can only assign households based on marital status (Chetty et al., 2014). The scope of fiscal households excludes individuals living in collective structures (retirements homes, religious communities, student accommodations, prisons, etc.) as well as those most in distress, who live in precarious housing (worker hostels, etc.) or are homeless.

B. PSID Validation Exercise

We use the Panel Study of Income Dynamics (PSID) to assess the extent to which OLS and two-sample two-stage least squares (TSTLS) estimates of rank-based intergenerational mobility measures differ from one another. Our sample and definition choices aim to be as close as possible to our main analysis setting while at the same time maximizing sample size. Note also that for this reason we use all of the PSID, rather than only the nationally-representative Survey Research Center (SRC) component. The main conclusions of our baseline results are robust to using only the SRC sample or to using various weighing schemes as shown in Section B.7.

B.1. Sample Definitions

Sample of Children. It consists of individuals who are (i) born between 1963 and 1988, (ii) observed as children in a family unit at least once, and (iii) observed at least once as reference person or partner in a family unit between 30 and 50 years old. Restriction (i) enables us to identify parents, while restriction (ii) enables us to observe children's incomes. The final sample contains 5,655 children.³²

Sample of Parents. Following Chetty et al. (2014), for each child, we define parent(s) as the reference person and partner of the family unit in which the child is first observed.³³ We then follow these individuals' incomes over time.³⁴ As Chetty et al. (2014), for simplicity, we fix each child's parent assignment regardless of any potential subsequent changes to the child's family unit reference person and partner. The final sample contains 5,785 (unique) parents.

B.2. Variable Definitions

All income variables are measured in 2019 dollars, adjusting for inflation using the consumer price index (CPI-U). Following Lee and Solon (2009) and Mazumder (2016), we exclude income observations obtained by "major assignment". We opt for larger age ranges than in our main analysis (30-50 vs 35-45) to increase our sample size. However, our baseline results are robust to averaging over 35-45 as in the main analysis (see Appendix Table B.5).

Parent Income. We rely on two parent income definitions. First, as a benchmark, we measure parent income as total pretax income at the household level, which we label parent family income. Specifically, we define parent family income as the sum of taxable income of the family unit's reference person and partner, and total transfer income of the reference person and partner.³⁵

³²See Appendix Table B.2 for the sample size at each additional restriction.

³³90% of individuals born in 1963-1988 are first observed as children in a family unit prior to age 18.

³⁴Note that this differs from the following studies using the PSID: Lee and Solon (2009) (use the family taxable income in which the children find themselves between ages 15 and 17), Mazumder (2016) (uses the PSID's Family Identification Mapping System (FIMS) to identify fathers), Jerrim et al. (2016) (do not explain exactly how fathers are identified; to be precise, the authors write "[...] we only include sons whose father can be identified," (Jerrim et al., 2016, p.89)), and Bloise et al. (2021) (do not explain exactly how fathers are identified; to be precise, the authors write "we include only sons whose real fathers have at least five years of positive earnings [...]" (Bloise et al., 2021, p.650)).

³⁵The accuracy of the family's taxable income is missing in 1993-1996 and in 2001-2019. Total transfers

Taxable income is equal to the sum of reference person's labor income, the partner's labor income, income from assets, and net profit from farm or business. This measure enables us to obtain benchmark estimates that the TSTSLS estimation strategy is supposed to yield.

Second, since in TSTSLS strategies parent family income is rarely observed, we also define parent labor income as the sum of family unit's reference person and partner's individual labor incomes (money income from labor, including self-employment income).³⁶ This follows very closely the setting adopted in the main analysis.

For both parent family income and parent labor income, we average income values over 30 and 50 years old. Specifically, we take the sum of the average for the father and the average for the mother if both parents are observed, and take the average of the only observed parent otherwise.

Child Income. We define child income in the same way as parent family income, again averaging over income observations between 30 to 50 years old.

Adjustment for Household Size. When defining income variables we follow [Chetty et al. \(2014\)](#), and do not account for household size (i.e., whether there is also a partner in the family unit). This way of defining parent income mechanically hinders single-headed households, both parents and children.³⁷ We therefore show in Table B.5 results when dividing family income measures by the number of observed reference person and partner in that year.

Descriptive Statistics Appendix Table B.1 displays some descriptive statistics for our sample of parents and children. Parents' incomes are observed at a slightly older age (39) than that of our children (34). In both cases, incomes are measured sufficiently late in the lifecycle to limit lifecycle bias.

are missing in 1968 and 1969. Total transfers include aid to families with dependent children, supplemental security income, other welfare payments, social security payments, other retirement, pensions and annuities, unemployment pay, workmen's compensation, child support, help from relatives, and other transfer income.

³⁶The accuracy of the reported value for the reference person is missing in 1994-1996. Moreover, for partners, there was a small change in income definition in 1994: total labor income became total labor income excluding farm and business income.

³⁷Interestingly, this is an issue Raj Chetty alludes to in his conversation with Tyler Cowen in his 2017 [Conversations with Tyler](#) podcast episode. Indeed, Chetty noticed that daughters from affluent families in the Bay area have low *household* incomes but have very high *individual* incomes because they are significantly less likely to be married than if they had grown up somewhere else.

Table B.1: Descriptive Statistics

	N	Missing (%)	Mean	Std. Dev.	25th pctile	Median	75th pctile
Parents							
Family income (average 30-50 yrs old)	5,785	5.88	82,047	66,121	42,976	72,081	105,523
Number of family income observations	5,785	5.88	13	5	10	15	18
Mean age at family income obs.	5,785	5.88	39	3	38	39	40
Labor income (average 30-50 yrs old)	5,785	5.62	39,679	39,946	13,575	30,800	55,074
Number of labor income observations	5,785	5.62	14	5	10	15	19
Mean age at labor income obs.	5,785	5.62	39	3	38	39	40
Fraction single parents	20.19%						
Fraction female among single parents	92.21%						
Mother's age at child birth	3,135	0.00	25	5	21	25	29
Father age at child birth	2,650	0.00	28	6	24	28	32
Children							
Family income (average 30-50 yrs old)	5,655	3.02	80,539	75,072	34,936	64,092	104,517
Number of family income observations	5,655	3.02	5	3	2	4	8
Mean age at family income obs.	5,655	3.02	34	3	32	34	38
Fraction female	53.60%						

Notes: See Sections B.1 and B.2 for details on sample construction and income definitions. Missing income observations can also correspond to values obtained by 'major assignment'.

B.3. Benchmark Estimates

We first estimate the rank-rank correlation (RRC) and transition matrix using the family income definitions for both parents and children (results for the intergenerational income elasticity (IGE) are presented in Appendix Figure B.5). Recall that in the TSTSLS setting the parent income definition is parent labor income while we are actually interested in parent family income, a more comprehensive parent income measure. In theory, the extent to which the additional incomes included in parent family income relative to parent labor income generate large rank reversals is ambiguous. Moreover, TSTSLS estimates necessitate restricting the analysis to the sample of children for whom parents' characteristics are observed (e.g., education and/or occupation, etc.). Such restrictions could potentially induce some biases relative to the statistic one is actually interested in measuring.

National Results. Appendix Figure B.1 displays the benchmark RRC and transition matrix for the baseline parent and child income definitions. The baseline RRC is 0.504, compared to 0.34 found in Chetty et al. (2014). Such a high RRC likely reflects the fact that the PSID contains oversamples of disadvantaged families (see Appendix Figure B.7 for estimates obtained only on the Survey Research Center (SRC) component of the PSID). The benchmark transition matrix confirms this intuition. The share of children from the bottom 20% who reach the top 20% in adulthood is 4%, close to half the share found by Chetty et al. (2014) (7.5% for children born in 1980-1982). Persistence at the bottom and top are also very strong at roughly 45%.

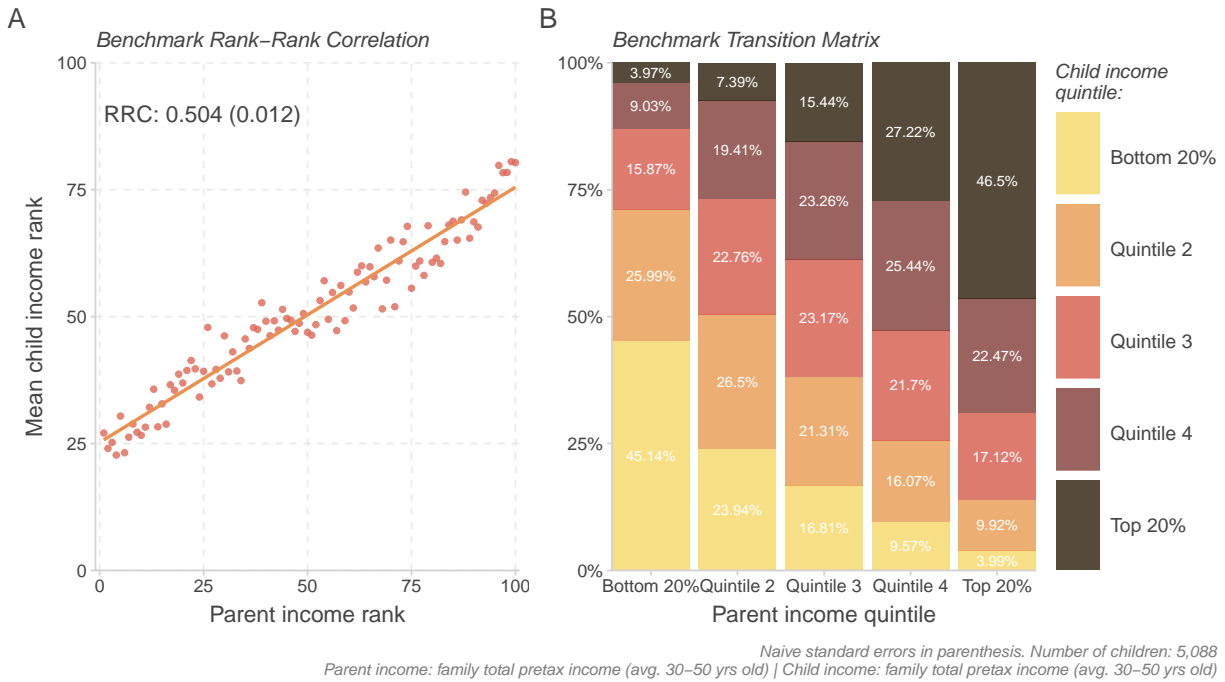


Figure B.1. Benchmark Rank-Rank Correlation and Transition Matrix

Notes: This figure presents the rank-rank correlation (panel A) and the transition matrix (panel B). It is computed on the Panel Study of Income Dynamics (PSID). The sample used is restricted to children born between 1963 and 1988 who are observed at least once as children in a family unit and at least once as a reference person or partner in a family unit over ages 30-50. Child income is the mean of family total income over ages 30-50. Parent income is the sum of father and mother mean family total income over ages 30-50. In panel A, the fitted line is a linear fit through the conditional expectation. We report coefficients and naive standard errors (in parenthesis) obtained from OLS regressions of child income rank on parent income rank with child cohort fixed effects, on the microdata for the full sample.

Subnational Results. Due to sample size constraints we explore geographic heterogeneity in intergenerational by Census Region (Northeast, Midwest, South, and West). Specifically, we define a child’s Census Region as the most common region of residence until age 18 (included). Appendix Figure B.2 displays the benchmark RRC and absolute upward mobility (AUM) estimates by Census Region. AUM is defined as in [Chetty et al. \(2014\)](#) as the expected income rank for children at the 25th percentile of the parent income distribution.




Figure B.2. Benchmark Rank-Rank Correlation and Absolute Upward Mobility by Census Region

Notes: This figure presents Census Region-level estimates of the rank-rank correlation (RRC) and absolute upward mobility (AUM). To compute local estimates, individuals are assigned to their most common Census Region of residence until age 18 (included). See Appendix Figure B.1’s notes for details on data, sample and income definitions.

B.4. OLS vs. TSTSLs Comparison

We now turn to the comparison between estimates obtained with OLS and those obtained with TSTSLs. The PSID enables us to compare estimates of intergenerational mobility we obtain when observing parents’ incomes and when predicting them using observable characteristics such as education and occupation. Since in the main analysis and in virtually all TSTSLs studies only parents’ labor incomes or wages are observed, we define parents’ income as individual labor income, while keeping in mind the benchmark estimates presented in the previous section. We follow the main analysis’ definitions as closely as possible. We proceed in the following way.

Parent Income Prediction

Let Z denote a set of characteristics observed for parents. We can express their labor incomes y as $y_i = \beta Z_i + \epsilon_i$. We estimate this first-stage equation by OLS on our sample of parents, and predict out of sample using a 5-fold cross-validation approach. Specifically, we split the sample of parents in five random subsamples of equal size, and for each subsample we predict income using the first-stage estimated on the remaining four subsamples. As such all predicted incomes are conceptually made from a random sample of parents taken from the same population. We see these out-of-sample predictions as imitating very closely settings in which researchers do not

observe the actual parents' incomes but observe the incomes of other parents taken from the same population (i.e., with children born in the same years).

We define parent income y as log mean (individual) labor income over ages 30 to 50. Once we have predicted labor incomes for children's father and/or mother, we compute a measure of labor income at the household level as the sum of father and mother predicted labor incomes if we have identified two parents, and predicted labor income of the only parent otherwise.³⁸ We display parents' (out-of-sample) predicted labor incomes against observed labor incomes in Appendix Section B.7.

For our baseline results, we define Z in the most similar way as possible as to our paper. Specifically, Z includes (i) education (7 categories; highest years of school completed), (ii) 3-digit occupation (334 cat.; most common occupation, including inactivity status, between 30 and 50 years old), (iii) demographic characteristics (birth cohort, race (5 cat.; most recent observation)), and (iv) state fixed effects (most common state of residence between 30 and 50 years old). The precise details of the construction of each of these variables are described in Appendix Section B.6. This set of predictors departs from the ones used in the main analysis because (i) we were unable to find a cross-walk between the 3-digit classification and a 2-digit classification, (ii) nationality is not available in the PSID, and (iii) country of birth is not available in the PSID. We replaced these variables with race. In Appendix Table B.4 we present results when incrementally including these predictors and find that the TSTSLS bias stabilizes once occupation is included.

National Results

Appendix Figure B.3 presents the main results from our validation exercise. Our TSTSLS estimate of the RRC is 0.459. On the exact same sample the OLS estimate is 0.476. Our benchmark RRC from the previous section was 0.504. The TSTSLS estimate is therefore roughly 4% smaller than the OLS estimate on the same sample, and 9% smaller than the benchmark OLS estimate (defining parent income as parent family income). These differences are quite small relative to the large differences in RRC estimates observed across countries (as well as within country across studies). Moreover, and importantly, the TSTSLS estimate appears to *understate* persistence, suggesting it provides a lower bound for intergenerational persistence.

The TSTSLS estimates for the transition matrix also appear to represent upper bounds on intergenerational (upward) mobility. The P(Top 20% — Bot. 20%) is roughly 6% in the TSTSLS case and 4% in the OLS case (4% as well in the benchmark), P(Bot. 20% — Bot. 20%) is 40% vs. 44% (45%) and the P(Top 20% — Top 20%) is 44% vs. 46% (47%). In Appendix B.7 we show that the TSTSLS bias of the RRC is largely unaffected by the number of parent income observations used. Moreover, Table B.5 shows our results are qualitatively robust to (i) using nationally-representative Survey Research Center (SRC) sample of the PSID, (ii) computing parent and child incomes over ages 35-45 as in the main analysis, (iii) dropping income observations equal to zero when computing parent and child incomes, and (iv) accounting for household size in the income definitions (additional details in Section B.7). Moreover, Table B.6 shows that using the longitudinal or cross-sectional weights moderately increases the TSTSLS RRC downward bias.

³⁸Results when dividing by the number of parents are presented in Appendix Table B.5 (col. (5)).

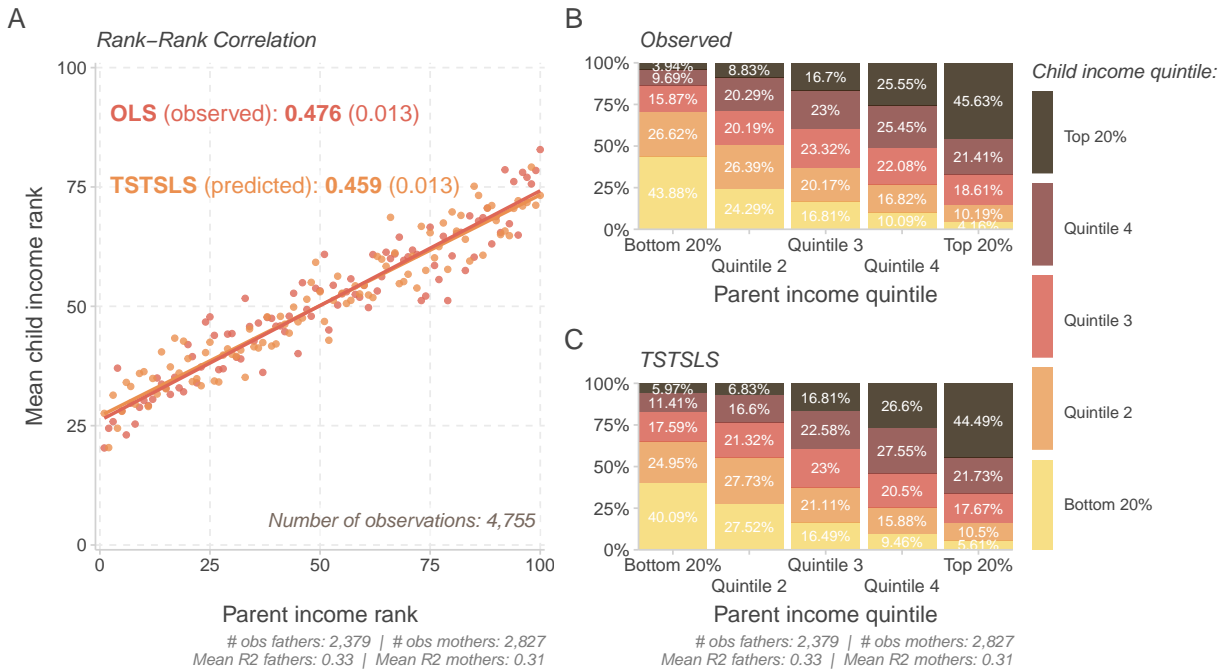


Figure B.3. OLS vs. TSTSLs RRC and Transition Matrix

Notes: This figure presents the rank-rank correlation (panel A) and the transition matrix (panels B and C) obtained when parent income is observed (OLS/observed) and when it is predicted using two-sample two-stage least squares (TSTSLs). It is computed on the Panel Study of Income Dynamics (PSID). The sample used is restricted to children born between 1963 and 1988 who are observed at least once as children in a family unit and at least once as a reference person or partner in a family unit over ages 30-50. Child income is the mean of family total income over ages 30-50. Parent income is the sum of father and mother (predicted) mean labor income over ages 30-50. For TSTSLs estimates, parent income is predicted separately for males and females using an OLS model including education (7 cat.; highest years of school completed), 3-digit occupation (334 cat.; most common occupation (incl. inactivity status) between 30 and 50 years old), demographic characteristics in 1990 (birth cohort and race (5 cat.; most recent observation) and state fixed effects (most common state of residence between 30 and 50 years old). In panel A, the fitted line is a linear fit through the microdata. We report coefficients and naive standard errors (in parenthesis) obtained from OLS regressions of child income rank on parent income rank with child cohort fixed effects, on the microdata for the full sample.

Regional Results

Appendix Figure B.4 shows the results obtained by Census Region. The RRC obtained by TSTSLs is remarkably similar to that obtained by OLS, with a slight underestimation for the Northeast and West regions. The same applies to the AUM which again is very similar in the TSTSLs setting relative to the OLS case (and the benchmarks). Compared to the benchmark estimates presented in the previous section, differences in RRCs are a bit larger but the rank-ordering of regions is preserved.

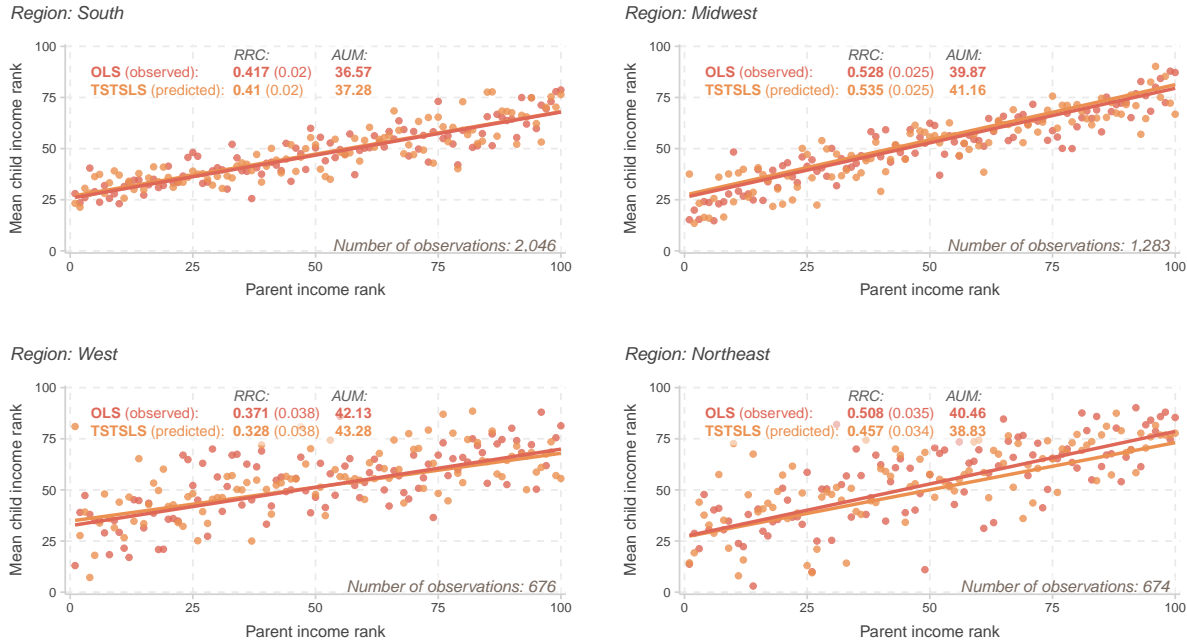


Figure B.4. OLS vs. TSTSLs RRC and AUM

Notes: This figure presents Census Region-level estimates of the rank-rank correlation (RRC) and absolute upward mobility (AUM). To compute local estimates, individuals are assigned to their most common Census Region of residence until age 18 (included). See Appendix Figure B.3's notes for details on data, sample and income definitions.

B.5. Discussion

Overall, the results presented in this analysis suggest that using TSTSLs for rank-based measures of intergenerational mobility leads to reasonably close estimates relative to OLS estimates, both at the national and subnational levels. Specifically TSTSLs estimates appear to slightly underestimate intergenerational persistence, from 4% to 10% depending on the set of predictors (see Appendix Section B.7 for all results when varying the set of first-stage predictors). Moreover, they seem to represent lower bounds for intergenerational persistence (i.e., upper bounds for mobility). In Appendix Table B.5, we show these findings are also robust to dropping income observations equal to 0, as well as to accounting for the number of reference person and partner when defining incomes for children and parents.

B.6. Details on Samples and Variable Definitions

Sample Construction Details

Table B.2: Sample Size at Each Restriction

	# obs.	%
Raw sample	82,573	100
+ born 1963-1988	30,186	36.56
+ observed at least once as child in a family unit	18,612	61.66
+ observed at least once as head/spouse 30-50	5,655	30.38
+ at least one family total income observation 30-50	5,484	96.98
+ at least one observation for parent total income observation 30-50	5,088	92.78

Notes: child and parent income observations exclude those obtained by "major assignment".

Details on Variable Constructions

Age: since prior to 1983, only age (rather than birth year) was reported, we use the following rule to obtain individuals' birth year: (i) if at least 1 birth year value: most common value; (ii) otherwise: most common value obtained from year - age (by definition this will equal birth year or birth year + 1).

Parent education: maximum grade completed over all observations, and classified following [Jerim et al. \(2016\)](#) / PSID classification of grades into education levels.³⁹

Categories: Grades 1-5, Grades 6-8, Grades 9-11, Grade 12 (HS completion), Some college / associate degree (grades 13-15), College degree (grade 16), Advanced college degree (grade 17).

Parent occupation: most common 3-digit occupation (1970 classification) or detailed inactivity status between 30 and 50 years old.⁴⁰ Occupation variables with a consistent classification are available for all individuals between 1981 and 2001, and are only available for a selected sample of PSID heads and wives/"wives"⁴¹ between 1968 and 1980. In order to prevent bias from focusing only on employed parents⁴², we use information from employment status variables from 1981 onwards.⁴³

³⁹Note the grade completed variable is missing for 1969.

⁴⁰In cases where an individual has several most common occupations, we assign the one for which the individual is the oldest on average, and choose one at random if average age is the same.

⁴¹Criteria: (i) original sample Heads and Wives/"Wives still living by 1992 who reported main jobs in at least three waves during the period 1968-1992, with at least one of those reports prior to 1980; and (ii) additionally, original sample Heads and Wives/"Wives" who had reported at least one main job between 1968 and 1980 but were known to have died by 1992. Those who were still living but had reported only one or two jobs during the period of interest were excluded, as were all nonsample Heads and Wives/"Wives".

⁴²By definition, occupations are only available for employed individuals.

⁴³Employment status is only available for heads between 1968 and 1978; from 1979 onwards, it is available for heads and wives/"wives". To prevent any bias, employment status is used only after 1980, i.e., when occupation is not restricted to a selected sample.

Categories: 441 3-digit occupations + 5 detailed inactivity status (Unemployed, Housewife, Student, Retired/Permanently disabled, Other).

Parent race: most recent race observation.

Categories: White, African American, Asian/Pacific Islander, Native American, Other.

Parent region: most common state between 30 and 50 years old.

Child region: most common state between 0 and 18 years old.

Table B.3: Detail of Variables Used in Each Year

Year	Int. # (ind. file)	Int. # (fam. files)	Sequence #	Rel. to ref. person	Occupation			Race			Employment status			Labor income			Accu. tax. inc.	Total transfers
					Age	Birth year	Education	Ref. person	Partner	Partner	State of res.	Ref. person	Accu.	Partner	Accu.	Partner		
1968	ER30001	V3	-	ER30003	ER30004	ER30010	V197-A	V243-A	V181	V93	V74	V85	V76	V86	V87	-	-	
1969	ER30020	V442	ER30021	ER30022	ER30023	ER30044	V60-A	V609-A	V801	V537	V514	V515	V516	V517	V518	-	-	
1970	ER30043	V1102	ER30044	ER30045	ER30046	ER30078	V1279-A	V1867-A	V1901	V1803	V1196	V1197	V1198	V1199	V1200	V1206	V1220	
1971	ER30068	V1802	ER30068	ER30069	ER30070	ER30076	V1984-A	V2074-A	V2202	V1803	V1897	V1898	V1899	V1900	V1906	V1907	V1922	V1922
1972	ER30091	V2402	ER30092	ER30093	ER30094	ER30100	V2582-A	V2672-A	V2828	V2403	V2498	V2499	V2500	V2501	V2508	V2523	V2523	V2523
1973	ER30117	V3002	ER30118	ER30119	ER30120	ER30126	V3115-A	V3183-A	V3300	V3003	V3051	V3047	V3052	V3053	V3060	V3076	V3076	V3076
1974	ER30138	V3402	ER30139	ER30140	ER30141	ER30162	V3968-A	V3601-A	V3720	V3403	V3463	V3459	V3465	V3466	V3473	V3488	V3488	V3488
1975	ER30160	V3802	ER30161	ER30162	ER30163	ER30169	V4459-A	V4055-A	V4204	V3803	V3869	V3859	V3865	V3866	V3872	V3889	V3889	V3889
1976	ER30188	V4302	ER30189	ER30190	ER30191	ER30197	V5374-A	V4605-A	V3996	V4303	V5031	V4379	V4386	V4387	V4387	V4404	V4404	V4404
1977	ER30217	V5202	ER30218	ER30219	ER30220	ER30226	V5873-A	V5807-A	V5662	V5203	V5627	V5284	V5289	V5290	V5297	V5316	V5316	V5316
1978	ER30246	V5702	ER30247	ER30248	ER30249	ER30255	V6497-A	V6039-A	V6209	V5703	V6174	V5783	V5788	V5789	V5797	V5815	V5815	V5815
1979	ER30283	V6302	ER30284	ER30285	ER30286	ER30296	V7100-A	V6996-A	V6802	V6303	V6767	V6392	V6398	V6399	V6406	V6426	V6426	V6426
1980	ER30313	V6902	ER30314	ER30315	ER30316	ER30326	V7329	V7198-A	V7447	V6903	V7413	V6982	V6988	V6989	V6997	V7016	V7016	V7016
1981	ER30343	V7502	ER30344	ER30345	ER30346	ER30356	V7712	V7885	V8099	V7503	V8066	V7574	V7580	V7581	V7589	V7608	V7608	V7608
1982	ER30373	V8202	ER30374	ER30375	ER30376	ER30384	V8380	V8544	V8723	V8203	V8690	V8266	V8273	V8274	V8282	V8301	V8301	V8301
1983	ER30399	V8802	ER30400	ER30401	ER30402	ER30413	V9011	V9194	V9408	V8803	V9376	V8874	V8881	V8882	V8890	V8909	V8909	V8909
1984	ER30429	V10002	ER30430	ER30431	ER30432	ER30443	V10678	V10678	V11055	V10003	V10223	V10257	V10263	V10264	V10277	V10305	V10305	V10305
1985	ER30463	V11102	ER30464	ER30465	ER30466	ER30478	V11651	V12014	V11938	V11103	V12295	V11398	V11404	V11405	V11419	V11418	V11418	V11418
1986	ER30498	V12502	ER30499	ER30500	ER30501	ER30513	V13054	V13233	V13565	V13500	V12803	V12797	V12803	V12804	V12818	V12868	V12868	V12868
1987	ER30535	V13702	ER30536	ER30537	ER30538	ER30549	V14154	V14329	V14671	V14547	V13703	V13899	V13905	V13906	V13920	V13970	V13970	V13970
1988	ER30570	V14802	ER30571	ER30572	ER30573	ER30584	V15162	V15464	V16086	V16021	V14803	V14914	V14920	V14921	V14934	V14985	V14985	V14985
1989	ER30606	V16802	ER30607	ER30608	ER30609	ER30620	V16663	V16982	V17483	V17418	V16303	V16414	V16420	V16421	V16435	V16485	V16485	V16485
1990	ER30642	V17702	ER30643	ER30644	ER30645	ER30657	V18101	V18403	V18834	V17703	V18878	V17830	V17856	V17857	V17851	V17901	V17901	V17901
1991	ER30689	V19002	ER30690	ER30691	ER30692	ER30703	V19401	V19703	V20114	V20049	V20178	V19130	V19136	V19137	V19150	V19201	V19201	V19201
1992	ER30733	V20502	ER30734	ER30735	ER30736	ER30748	V20701	V21003	V21420	V21355	V20303	V20430	V20436	V20437	V20450	V20501	V20501	V20501
1993	ER30806	V21602	ER30807	ER30808	ER30809	ER30820	V22456	V22809	V23276	V23212	V21603	V23323	V23324	V23324	V23366	V23366	V23366	V23366
1994	ER33101	ER5002	ER33102	ER33103	ER33104	ER33115	ER4017	ER4048	ER3944	ER3883	ER4156	ER4140	ER4144	ER4145	ER4147	ER4147	ER4147	ER4147
1995	ER33201	ER7002	ER33202	ER33203	ER33204	ER33215	ER6857	ER6888	ER6814	ER6753	ER6996	ER6980	ER6984	ER6986	ER6987	ER6987	ER6987	ER6987
1996	ER33301	ER7002	ER33302	ER33303	ER33304	ER33315	ER9108	ER9139	ER9060	ER8999	ER9247	ER9231	ER9236	ER9237	ER9238	ER9238	ER9238	ER9238
1997	ER33401	ER10002	ER33402	ER33403	ER33404	ER33415	ER12085	ER12116	ER11848	ER11760	ER12221	ER12080	ER12081	ER12082	ER12071	ER12071	ER12071	ER12071
1999	ER33501	ER13002	ER33502	ER33503	ER33504	ER33515	ER13215	ER13727	ER15928	ER15836	ER13004	ER16463	ER16465	ER16452	ER16453	ER16454	ER16454	ER16454
2001	ER33601	ER17002	ER33602	ER33603	ER33604	ER33616	ER17226	ER17796	ER19989	ER19987	ER17004	ER20445	ER20447	ER20448	ER20450	ER20450	ER20450	ER20450
2003	ER33701	ER21002	ER33702	ER33703	ER33704	ER33716	-	-	ER23426	ER23334	ER21003	ER21116	ER21136	ER21100	ER21001	ER21001	ER21001	ER21001
2005	ER33801	ER30002	ER33802	ER33803	ER33804	ER33815	-	-	ER27363	ER27297	ER25003	ER27953	ER27944	ER27953	ER28002	ER28002	ER28002	ER28002
2007	ER33901	ER36002	ER33902	ER33903	ER33904	ER33917	-	-	ER40585	ER40472	ER36003	ER40921	ER40933	ER40934	ER40992	ER40992	ER40992	ER40992
2009	ER34001	ER42002	ER34002	ER34003	ER34004	ER34020	-	-	ER46543	ER46449	ER42003	ER46829	ER46841	ER46851	ER46900	ER46900	ER46900	ER46900
2011	ER34101	ER47302	ER34102	ER34103	ER34104	ER34119	-	-	ER51904	ER51810	ER47303	ER48242	ER48249	ER48250	ER48308	ER48308	ER48308	ER48308
2013	ER34201	ER53002	ER34202	ER34203	ER34204	ER34230	-	-	ER57549	ER57549	ER53003	ER58038	ER58050	ER58060	ER58117	ER58117	ER58117	ER58117
2015	ER34301	ER60002	ER34302	ER34303	ER34304	ER34349	-	-	ER64671	ER64671	ER60003	ER65216	ER60995	ER60524	ER60525	ER60525	ER60525	ER60525
2017	ER34501	ER66002	ER34502	ER34503	ER34504	ER34548	-	-	ER70882	ER70744	ER66003	ER71293	ER70747	ER71330	ER71331	ER71331	ER71331	ER71331
2019	ER34701	ER72002	ER34702	ER34703	ER34704	ER34752	-	-	ER76897	ER76752	ER72003	ER77315	ER73070	ER77343	ER77352	ER77352	ER77352	ER77352

Notes: Int. = interview. Ind. = individual. Fam. = family. Rel. = relation. Res. = residence. Accu. = accuracy. Tax. inc. = taxable income.

B.7. Additional Results

All Benchmark Estimates

Parent income definition	Individual labor income (sum) (30–50)	0.316 (0.011)	0.363 (0.012)	0.313 (0.012)	0.359 (0.013)	0.282 (0.009)	0.332 (0.011)	0.295 (0.013)
	Individual labor income (mean) (30–50)	0.359 (0.012)	0.405 (0.013)	0.355 (0.014)	0.4 (0.015)	0.32 (0.011)	0.369 (0.012)	0.338 (0.015)
	Family total income (30–50)	0.531 (0.016)	0.607 (0.017)	0.547 (0.018)	0.622 (0.019)	0.503 (0.014)	0.584 (0.015)	0.487 (0.019)
	Family total income (div. number of adults) (30–50)	0.612 (0.018)	0.682 (0.02)	0.628 (0.021)	0.698 (0.023)	0.584 (0.016)	0.659 (0.018)	0.562 (0.023)
	Family taxable income (30–50)	0.348 (0.011)	0.4 (0.012)	0.353 (0.013)	0.406 (0.014)	0.319 (0.01)	0.375 (0.011)	0.323 (0.014)
	Family taxable income (div. number of adults) (30–50)	0.38 (0.012)	0.431 (0.013)	0.384 (0.014)	0.435 (0.015)	0.348 (0.011)	0.403 (0.012)	0.355 (0.015)
	Family labor income (30–50)	0.369 (0.012)	0.425 (0.013)	0.373 (0.013)	0.429 (0.014)	0.333 (0.01)	0.392 (0.012)	0.343 (0.015)
	Family labor income (div. number of adults) (30–50)	0.405 (0.013)	0.458 (0.014)	0.408 (0.015)	0.461 (0.016)	0.364 (0.011)	0.421 (0.013)	0.377 (0.016)
		Family labor income (div. number of adults) (30–50)	Family taxable income (div. number of adults) (30–50)	Family total income (div. number of adults) (30–50)	Family total income (30–50)	Individual labor income (30–50)		
		Child income definition						

Naive standard errors in parenthesis. Number of children varies by income definition since the number of negative or zero incomes varies.

Figure B.5. Benchmark IGEs for All Income Definitions

Parent income definition	Individual labor income (sum) (30–50)	0.465 (0.012)	0.47 (0.012)	0.47 (0.012)	0.474 (0.012)	0.476 (0.012)	0.48 (0.012)	0.371 (0.013)
	Individual labor income (mean) (30–50)	0.458 (0.012)	0.453 (0.013)	0.463 (0.012)	0.456 (0.013)	0.47 (0.012)	0.462 (0.012)	0.37 (0.013)
	Family total income (30–50)	0.487 (0.012)	0.491 (0.012)	0.495 (0.012)	0.497 (0.012)	0.503 (0.012)	0.504 (0.012)	0.386 (0.013)
	Family total income (div. number of adults) (30–50)	0.479 (0.012)	0.471 (0.012)	0.486 (0.012)	0.476 (0.012)	0.495 (0.012)	0.483 (0.012)	0.38 (0.013)
	Family taxable income (30–50)	0.491 (0.012)	0.496 (0.012)	0.499 (0.012)	0.502 (0.012)	0.506 (0.012)	0.509 (0.012)	0.39 (0.013)
	Family taxable income (div. number of adults) (30–50)	0.485 (0.012)	0.481 (0.012)	0.493 (0.012)	0.487 (0.012)	0.501 (0.012)	0.493 (0.012)	0.386 (0.013)
	Family labor income (30–50)	0.48 (0.012)	0.483 (0.012)	0.484 (0.012)	0.486 (0.012)	0.491 (0.012)	0.492 (0.012)	0.385 (0.013)
	Family labor income (div. number of adults) (30–50)	0.472 (0.012)	0.465 (0.012)	0.475 (0.012)	0.467 (0.012)	0.482 (0.012)	0.473 (0.012)	0.38 (0.013)
		Family labor income (div. number of adults) (30–50)	Family taxable income (div. number of adults) (30–50)	Family total income (div. number of adults) (30–50)	Family total income (30–50)	Individual labor income (30–50)		
		Child income definition						

Naive standard errors in parenthesis. Number of children: 5,088.

Figure B.6. Benchmark RRCs for All Income Definitions

Only SRC Sample




Figure B.7. Benchmark Rank-Rank Correlation and Transition Matrix

Notes: This figure presents the rank-rank correlation (panel A) and the transition matrix (panel B), computed on the Panel Study of Income Dynamics (PSID)’s representative Survey Research Center sample. See Appendix Figure B.1’s notes for details on data, sample and income definitions.

Baseline Predictions



Figure B.8. Observed vs. (out-of-sample) predicted *individual* labor income

Notes: This figure presents observed *individual* log labor income and out-of-sample predicted *individual* log labor income for fathers and mothers depending on variables used in the first-stage prediction. The red line corresponds to the 45 degree line. See Appendix Figure B.3's notes for details on data, sample and income definitions.




Figure B.9. Observed vs. (out-of-sample) predicted *individual* labor income rank

Notes: This figure presents the conditional expectation of out-of-sample predicted *individual* labor income rank, as a function of observed *individual* labor income rank, for fathers and mothers, depending on variables used in the first-stage prediction. The red line corresponds to the 45 degree line, while the red shaded area corresponds to the interquartile range. See Appendix Figure B.3's notes for details on data, sample and income definitions.

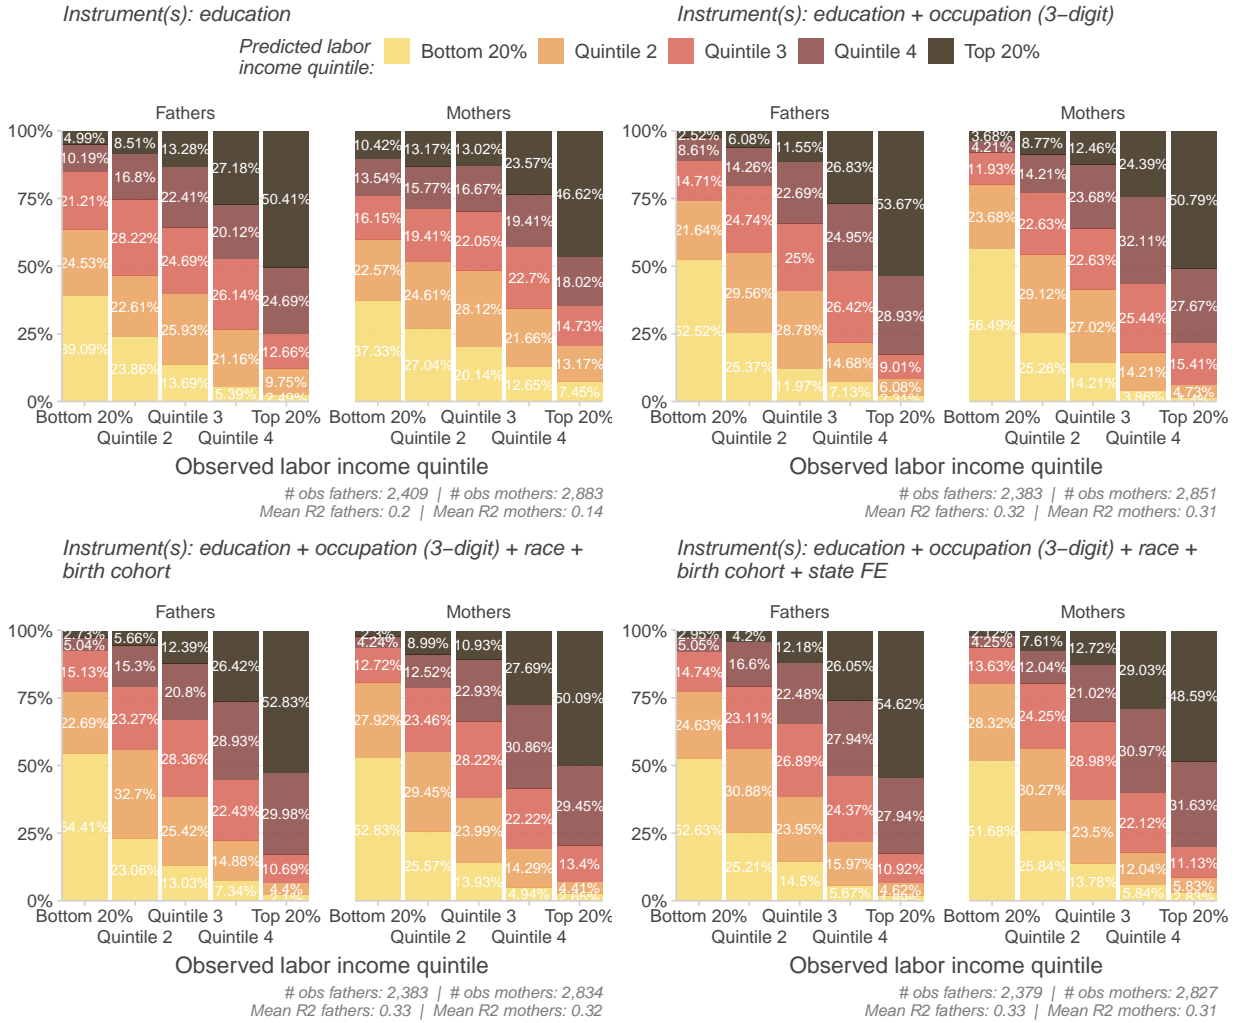


Figure B.10. Observed vs. (out-of-sample) predicted *individual* labor income quintile

Notes: This figure presents the quintile-by-quintile out-of-sample predicted *individual* labor income quintile by observed *individual* labor income quintile, for fathers and mothers, depending on variables used in the first-stage prediction. Each cell documents the share of out-of-sample labor income predictions belonging to the quintile indicated by the color legend among observed labor incomes falling in the quintile indicated on the x-axis. They are computed separately for father and mother, and depending on variables used in the first-stage prediction. See Appendix Figure B.3's notes for details on data, sample and income definitions.



Figure B.11. Observed vs. (out-of-sample) predicted *family* labor income

Notes: This figure presents observed *family* log labor income and out-of-sample predicted *family* log labor income for fathers and mothers depending on variables used in the first-stage prediction. The red line corresponds to the 45 degree line. See Appendix Figure B.3's notes for details on data, sample and income definitions.

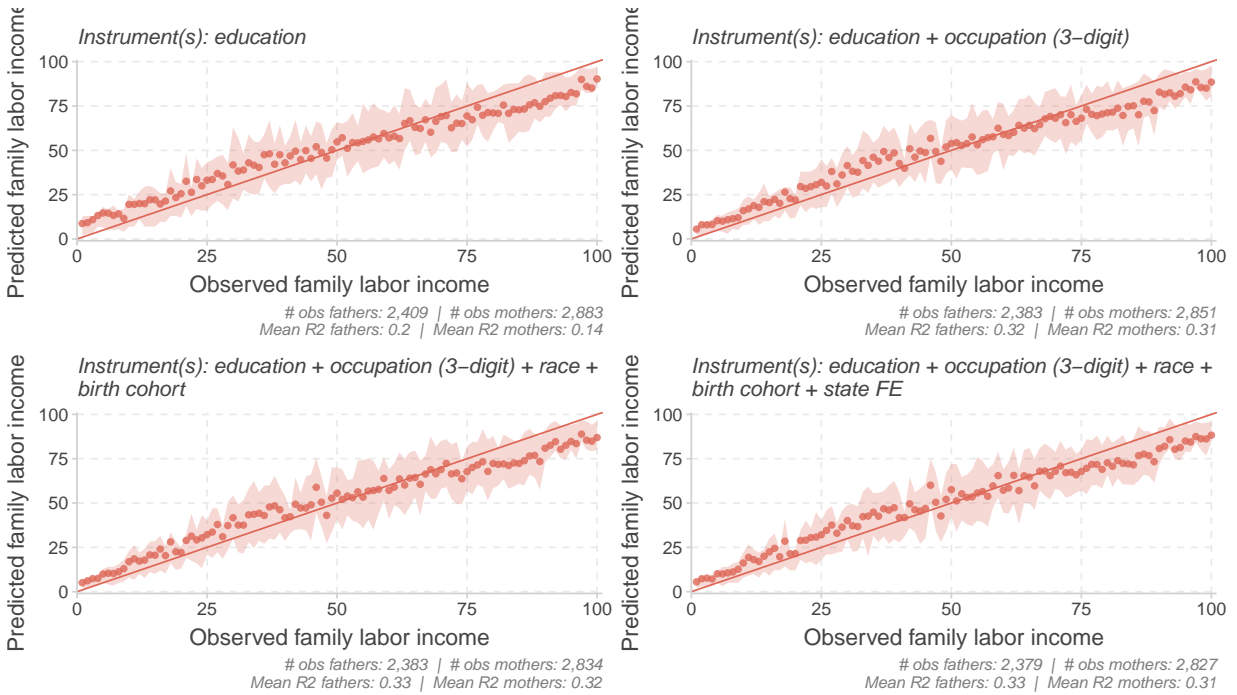


Figure B.12. Observed vs. (out-of-sample) predicted *family* labor income rank

Notes: This figure presents the conditional expectation of out-of-sample predicted *family* labor income rank, as a function of observed *family* labor income rank, for fathers and mothers, depending on variables used in the first-stage prediction. The red line corresponds to the 45 degree line, while the red shaded area corresponds to the interquartile range. See Appendix Figure B.3's notes for details on data, sample and income definitions.

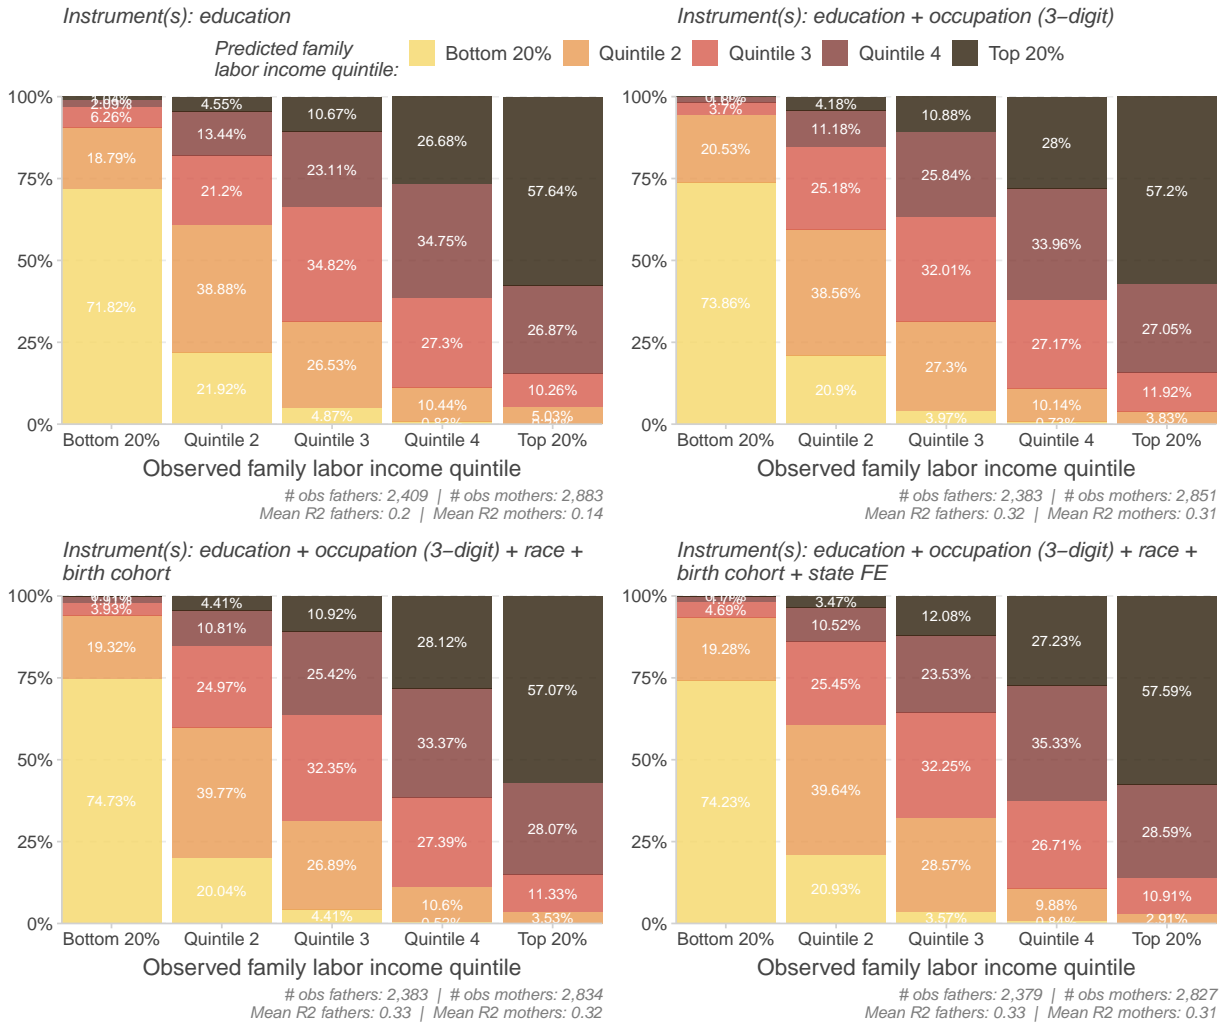


Figure B.13. Observed vs. (out-of-sample) predicted *family* labor income rank

Notes: This figure presents the quintile-by-quintile out-of-sample predicted *family* labor income quintile by observed *family* labor income quintile, for fathers and mothers, depending on variables used in the first-stage prediction. Each cell documents the share of out-of-sample labor income predictions belonging to the quintile indicated by the color legend among observed labor incomes falling in the quintile indicated on the x-axis. They are computed separately for father and mother, and depending on variables used in the first-stage prediction. See Appendix Figure B.3’s notes for details on data, sample and income definitions.

Table B.4: Comparison for Different Sets of Predictors

	Education (1)	+ occupation (3-digit) (2)	+ race + birth cohort (3)	+ state FE (4)
<i>Panel A. Intergenerational Elasticity (IGE)</i>				
Observed parent income (OLS)	0.335 (0.011)	0.334 (0.011)	0.334 (0.011)	0.334 (0.011)
Predicted parent income (TSTSLS)	0.464 (0.017)	0.431 (0.014)	0.449 (0.014)	0.445 (0.014)
Percentage diff. TSTSLS vs OLS	-27.76%	-22.57%	-25.73%	-24.82%
Number of observations	4,805	4,755	4,737	4,730
<i>Panel B. Rank-Rank Correlation (RRC)</i>				
Observed parent income (OLS)	0.476 (0.013)	0.475 (0.013)	0.476 (0.013)	0.476 (0.013)
Predicted parent income (TSTSLS)	0.43 (0.013)	0.453 (0.013)	0.461 (0.013)	0.459 (0.013)
Percentage diff. TSTSLS vs OLS	10.53%	4.9%	3.22%	3.85%
Number of observations	4,832	4,780	4,762	4,755
<i>Panel C. Transition Matrix</i>				
P(Bottom 20% — Bottom 20%) (OLS)	43.95%	43.7%	43.84%	43.88%
P(Bottom 20% — Bottom 20%) (TSTSLS)	37.83%	39.77%	40.66%	40.09%
P(Bottom 20% — Top 20%) (OLS)	4.51%	4.46%	4.26%	4.16%
P(Bottom 20% — Top 20%) (TSTSLS)	4.81%	5.18%	5.39%	5.61%
P(Top 20% — Bottom 20%) (OLS)	3.97%	3.92%	3.93%	3.94%
P(Top 20% — Bottom 20%) (TSTSLS)	5.22%	5.83%	5.63%	5.97%
P(Top 20% — Top 20%) (OLS)	45.54%	45.49%	45.53%	45.63%
P(Top 20% — Top 20%) (TSTSLS)	44.22%	43.32%	43.36%	44.49%

Notes:

Alternative Samples and Definitions

In Table B.5 we check the robustness of our baseline results to changes in estimation samples and income definitions. Specifically, we report results for the following changes: (i) using only the nationally-representative Survey Research Center (SRC) sample (col. 2), (ii) restricting the age range over which child and parent incomes are averaged to 35-45 years old in our main analysis, (iii) dropping parent and child income observations equal to zero when computing average incomes⁴⁴, (iv) accounting for household size when defining parent and child incomes (see discussion in B.2). Our baseline results are reported in column 1.

⁴⁴According to Mazumder (2016, p.101): "In the PSID, the household head is recorded as having zero labor income if their income was actually zero or if their labor income is missing, so one cannot cleanly distinguish true zeroes with labor income."

Table B.5: Robustness of Baseline Results

	Baseline Estimates (1)	Only SRC Sample (2)	35-45 Income Age Range (3)	Dropping Zero Inc. Obs. (4)	Accounting Household Size (5)
<i>Panel A. National - Intergenerational Elasticity (IGE)</i>					
Observed parent income (OLS)	0.334 (0.011)	0.369 (0.018)	0.363 (0.016)	0.414 (0.013)	0.324 (0.011)
Predicted parent income (TSTSLs)	0.445 (0.014)	0.418 (0.022)	0.475 (0.021)	0.53 (0.017)	0.485 (0.016)
Percentage diff. TSTSLs vs OLS	-24.82%	-11.72%	-23.53%	-21.9%	-33.21%
Number of observations	4,730	2,892	2,882	4,732	4,730
<i>Panel B. National - Rank-Rank Correlation (RRC)</i>					
Observed parent income (OLS)	0.476 (0.013)	0.409 (0.017)	0.464 (0.017)	0.47 (0.013)	0.466 (0.013)
Predicted parent income (TSTSLs)	0.459 (0.013)	0.364 (0.017)	0.448 (0.017)	0.463 (0.013)	0.435 (0.013)
Percentage diff. TSTSLs vs OLS	3.85%	12.38%	3.56%	1.7%	7.16%
Number of observations	4,755	2,903	2,903	4,732	4,755
<i>Panel C. Region: Midwest</i>					
RRC - OLS	0.528 (0.025)	0.417 (0.03)	0.506 (0.031)	0.523 (0.025)	0.509 (0.025)
AUM - OLS	39.87	43.94	38.49	39.65	37.19
RRC - TSTSLs	0.535 (0.025)	0.37 (0.032)	0.506 (0.032)	0.521 (0.025)	0.497 (0.026)
AUM - TSTSLs	41.16	46.85	40.29	41.43	39.01
RRC percentage diff. TSTSLs vs OLS	-1.22%	12.49%	0.02%	0.32%	2.52%
AUM percentage diff. TSTSLs vs OLS	-3.15%	-6.22%	-4.47%	-4.31%	-4.67%
Number of observations	1,283	980	834	1,277	1,283
<i>Panel D. Region: Northeast</i>					
RRC - OLS	0.508 (0.035)	0.429 (0.04)	0.457 (0.048)	0.503 (0.035)	0.52 (0.036)
AUM - OLS	40.46	41.91	46.3	40.23	44.23
RRC - TSTSLs	0.457 (0.034)	0.35 (0.039)	0.377 (0.045)	0.459 (0.033)	0.46 (0.035)
AUM - TSTSLs	38.83	41.49	46.9	37.85	42.34
RRC percentage diff. TSTSLs vs OLS	11.25%	22.4%	21.26%	9.6%	13.13%
AUM percentage diff. TSTSLs vs OLS	4.19%	1%	-1.28%	6.3%	4.44%
Number of observations	674	538	406	669	674
<i>Panel E. Region: South</i>					
RRC - OLS	0.417 (0.02)	0.398 (0.03)	0.423 (0.025)	0.414 (0.02)	0.413 (0.02)
AUM - OLS	36.57	36.19	35.4	37.17	36.71
RRC - TSTSLs	0.41 (0.02)	0.401 (0.03)	0.421 (0.026)	0.42 (0.021)	0.386 (0.02)
AUM - TSTSLs	37.28	35.01	35.34	37.47	38.12
RRC percentage diff. TSTSLs vs OLS	1.72%	-0.75%	0.47%	-1.5%	7.03%
AUM percentage diff. TSTSLs vs OLS	-1.92%	3.37%	0.15%	-0.8%	-3.71%
Number of observations	2,046	885	1,242	2,036	2,046
<i>Panel F. Region: West</i>					
RRC - OLS	0.371 (0.038)	0.299 (0.047)	0.321 (0.053)	0.357 (0.037)	0.353 (0.037)
AUM - OLS	42.13	40.03	45.57	42.11	42.68
RRC - TSTSLs	0.328 (0.038)	0.232 (0.046)	0.303 (0.052)	0.344 (0.038)	0.303 (0.038)
AUM - TSTSLs	43.28	43.43	46.18	42.98	43.66
RRC percentage diff. TSTSLs vs OLS	12.97%	28.89%	5.93%	3.8%	16.65%
AUM percentage diff. TSTSLs vs OLS	-2.66%	-7.84%	-1.31%	-2.02%	-2.23%
Number of observations	676	488	376	674	676




Figure B.14. OLS vs. TSTLS Estimates - Varying Number of Parent Income Observations

Notes: This figure presents the IGE, RRC and transition matrix cells obtained when parent income is observed (OLS) and when it is predicted using two-sample two-stage least squares (TSTLS), for different number of parent income observations. To control for the potential effect of lifecycle bias, we center parent incomes around age 40. Thus one observation means parent income is equal to income at age 40, two observations means parent income is equal to income averaged over age 39 and age 41, three observations means parent income is equal to income averaged over age 39 to age 41, etc. See Appendix Figure B.3's notes for details on data, sample and income definitions.

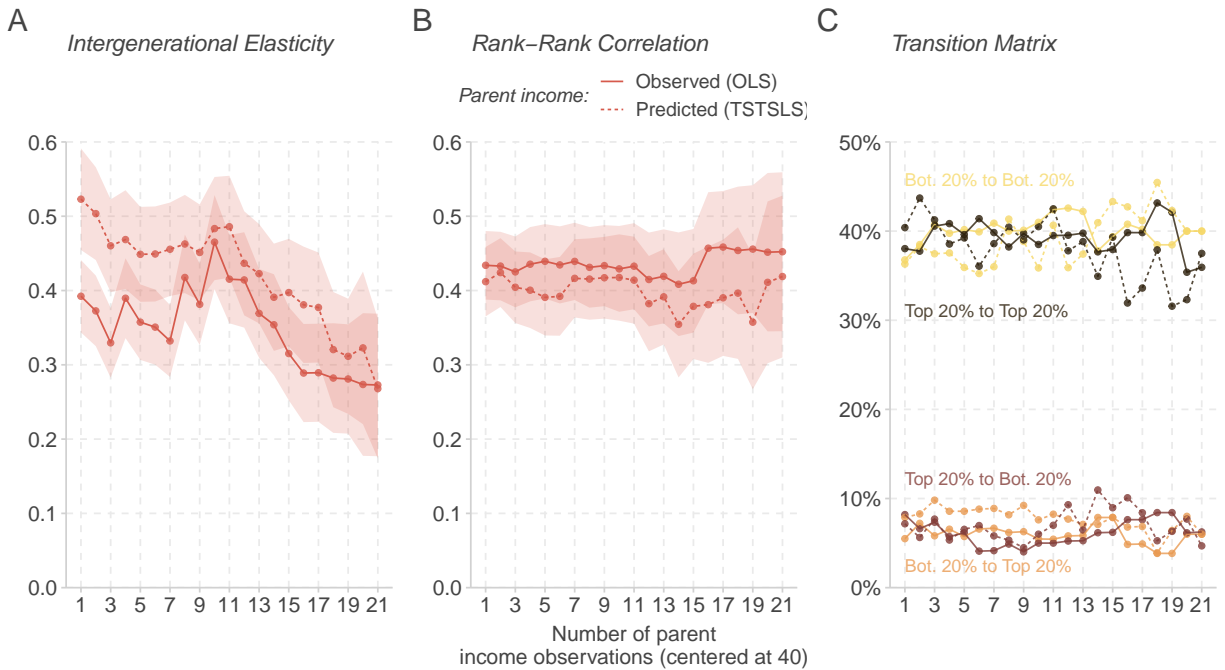


Figure B.15. OLS vs. TSTSLS Estimates - Varying Number of Parent Income Observations - Child Income Mean 37-43

Notes: This figure presents the IGE, RRC and transition matrix cells obtained when parent income is observed (OLS) and when it is predicted using two-sample two-stage least squares (TSTSLS), for different number of parent income observations and when child income is defined over ages 37-43. To control for the potential effect of lifecycle bias, we center parent incomes around age 40. Thus one observation means parent income is equal to income at age 40, two observations means parent income is equal to income averaged over age 39 and age 41, three observations means parent income is equal to income averaged over age 39 to age 41, etc. See Appendix Figure B.3's notes for details on data, sample and income definitions.

Sampling Weights

As is well-known, the PSID is not a nationally representative sample. In particular, the Survey of Economic Opportunity (SEO) component of the PSID oversamples low-income households but suffers from various sampling issues (see footnote 4 in [Lee and Solon \(2009\)](#)). In our baseline results, we opted to use all of the PSID because (i) our goal was to compare OLS to TSTSLS estimates rather than obtain the best OLS estimate, and (ii) the additional sample size allows us to compare OLS and TSTSLS estimates at the regional level. However, one may wish to know how our exercise performs for a nationally-representative sample. Table B.6 compares our baseline results with estimates obtained from four different specifications: (i) using only the PSID's nationally representative Survey Research Center (SRC) sample, (ii) using all of the PSID with three different kinds of weights, all measured in the child's last income observation year: (i) the family longitudinal weights, (ii) the individual longitudinal weights, and (iii) the individual cross-sectional weights (only available from 1997 onwards).

Overall, our baseline estimates have the smallest differences between TSTSLS and OLS. The OLS RRC is roughly 4% larger than the TSTSLS RRC in baseline, while it 12% when using only the SRC sample, 11% when using the family longitudinal weights, 9% when using individual longitudinal weights, and 12% when using the individual cross-sectional weights. Regarding the regional estimates, as with the baseline results, the relative difference between TSTSLS and OLS estimates largely reflect sample size: estimates for the Midwest and the South are quite close across specifications (a bit less so for the Midwest when using the SRC sample), while the differences become more pronounced for the Northeast and the West, for which the sample size is more limited. It should be noted that across regions and specifications, the TSTSLS estimates of the AUM are surprisingly close to their OLS counterparts.

Table B.6: Comparison between Baseline Results and Weighted Results

	<i>Weights in Last Child Income Observation Year</i>				
	Baseline	Only SRC	Family	Individual	Individual
	Estimates	Sample	Longitudinal	Longitudinal	Cross-Sectional
	(1)	(2)	(3)	(4)	(5)
<i>Panel A. National - Intergenerational Elasticity (IGE)</i>					
Observed parent income (OLS)	0.334 (0.011)	0.369 (0.018)	0.372 (0.012)	0.377 (0.012)	0.369 (0.013)
Predicted parent income (TSTSLs)	0.445 (0.014)	0.418 (0.022)	0.432 (0.014)	0.428 (0.014)	0.411 (0.015)
Percentage diff. TSTSLs vs OLS	-24.82%	-11.72%	-13.89%	-11.91%	-10.23%
Number of observations	4,730	2,892	4,730	4,730	4,588
<i>Panel B. National - Rank-Rank Correlation (RRC)</i>					
Observed parent income (OLS)	0.476 (0.013)	0.409 (0.017)	0.458 (0.013)	0.456 (0.013)	0.437 (0.014)
Predicted parent income (TSTSLs)	0.459 (0.013)	0.364 (0.017)	0.413 (0.013)	0.418 (0.013)	0.39 (0.014)
Percentage diff. TSTSLs vs OLS	3.85%	12.38%	10.91%	9.09%	12.07%
Number of observations	4,755	2,903	4,755	4,755	4,612
<i>Panel C. Region: Midwest</i>					
RRC - OLS	0.528 (0.025)	0.417 (0.03)	0.45 (0.026)	0.456 (0.026)	0.447 (0.027)
AUM - OLS	39.87	43.94	48.2	48.81	51.49
RRC - TSTSLs	0.535 (0.025)	0.37 (0.032)	0.427 (0.027)	0.433 (0.027)	0.444 (0.028)
AUM - TSTSLs	41.16	46.85	50.13	50.4	53.01
RRC percentage diff. TSTSLs vs OLS	-1.22%	12.49%	5.48%	5.25%	0.66%
AUM percentage diff. TSTSLs vs OLS	-3.15%	-6.22%	-3.85%	-3.15%	-2.87%
Number of observations	1,283	980	1,283	1,283	1,265
<i>Panel D. Region: Northeast</i>					
RRC - OLS	0.508 (0.035)	0.429 (0.04)	0.497 (0.036)	0.494 (0.036)	0.487 (0.037)
AUM - OLS	40.46	41.91	37.99	42.71	39.76
RRC - TSTSLs	0.457 (0.034)	0.35 (0.039)	0.443 (0.035)	0.441 (0.035)	0.42 (0.037)
AUM - TSTSLs	38.83	41.49	37.09	41.41	40.1
RRC percentage diff. TSTSLs vs OLS	11.25%	22.4%	12.24%	11.99%	15.77%
AUM percentage diff. TSTSLs vs OLS	4.19%	1%	2.42%	3.14%	-0.84%
Number of observations	674	538	674	674	650
<i>Panel E. Region: South</i>					
RRC - OLS	0.417 (0.02)	0.398 (0.03)	0.447 (0.019)	0.428 (0.019)	0.421 (0.02)
AUM - OLS	36.57	36.19	36.69	43.12	40.85
RRC - TSTSLs	0.41 (0.02)	0.401 (0.03)	0.437 (0.019)	0.423 (0.019)	0.406 (0.02)
AUM - TSTSLs	37.28	35.01	37.31	43.2	41.35
RRC percentage diff. TSTSLs vs OLS	1.72%	-0.75%	2.29%	1.25%	3.79%
AUM percentage diff. TSTSLs vs OLS	-1.92%	3.37%	-1.64%	-0.2%	-1.19%
Number of observations	2,046	885	2,046	2,046	1,970
<i>Panel F. Region: West</i>					
RRC - OLS	0.371 (0.038)	0.299 (0.047)	0.363 (0.038)	0.385 (0.039)	0.336 (0.04)
AUM - OLS	42.13	40.03	40.98	43.83	47.9
RRC - TSTSLs	0.328 (0.038)	0.232 (0.046)	0.264 (0.036)	0.301 (0.038)	0.231 (0.038)
AUM - TSTSLs	43.28	43.43	44.76	46.89	52.58
RRC percentage diff. TSTSLs vs OLS	12.97%	28.89%	37.34%	27.62%	45.79%
AUM percentage diff. TSTSLs vs OLS	-2.66%	-7.84%	-8.43%	-6.53%	-8.9%
Number of observations	676	488	676	676	651

C. Additional Robustness

This Appendix provides additional robustness checks to those presented in the body of the paper.

C.1. Sensitivity to Data Coverage

Civil Servants

We ensure our results are not affected by the fact that civil servants are only observed from 1988 onwards by estimating the first-stage regression computing synthetic parents' on post-1988 wages only, still restricting to when they are between 35 and 45 years old. Appendix Figure C.1 displays the results from this check. The results are largely unaffected.

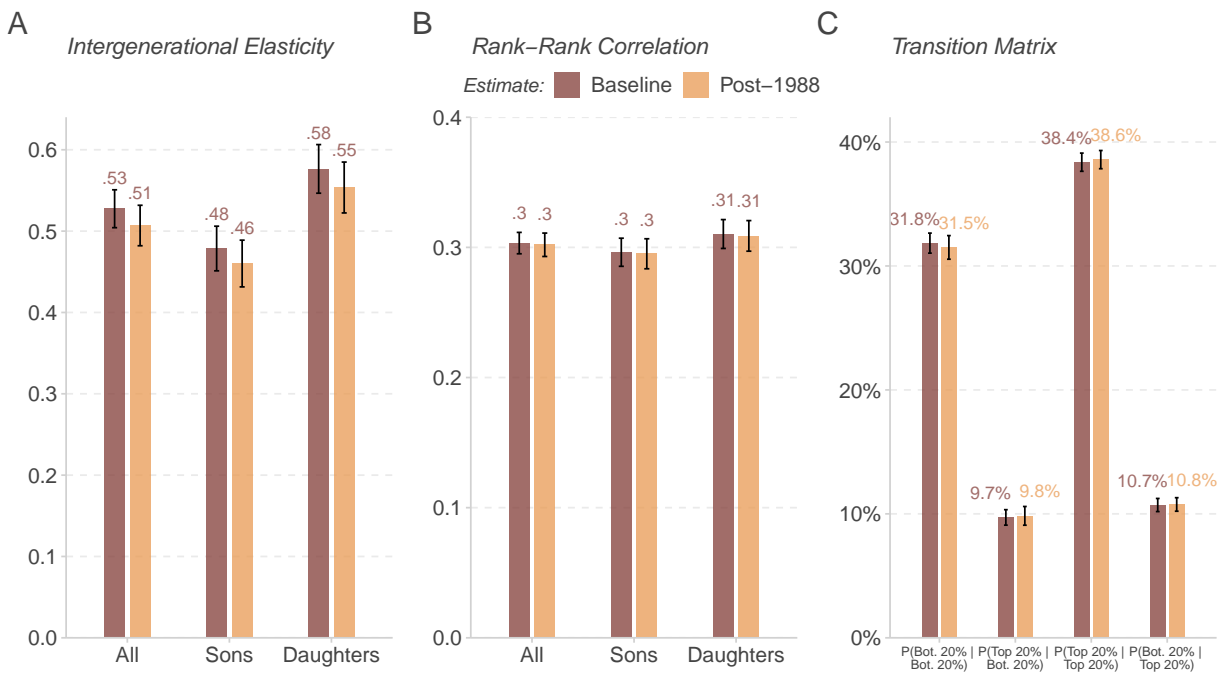


Figure C.1. Robustness of Baseline Estimates to Computing Synthetic Parent Incomes only on Post-1988 Data

Notes: This figure assesses the robustness of our baseline intergenerational income mobility estimates to computing synthetic parents' incomes only on post-1988 data. The All Employee Panel from which synthetic parents' wages are observed did not cover civil servants prior to 1988 (see Appendix Section A for details). The graph presents the baseline estimates (Baseline) to those obtained when synthetic parent incomes are defined as average wage between 35-45 using only post-1988 wages (Post-1988). Vertical lines represent the 95% bootstrapped confidence intervals. All results pertain to parent and child incomes being defined at the household level. The results for the transition matrix correspond to the sample pooling sons and daughters. See Section 3 for details on data, sample and income definitions.

Comparison with Population Statistics

Since the sample selection of the EDP is (virtually) random (individuals born on the first four days of October), we can have a good idea of how our baseline sample compares with the French population by comparing its average characteristics to those of the completely unrestricted EDP sample for the same birth cohorts (1972-1981).

To obtain characteristics on parents (other than from the 1990 census), we rely on individuals' birth-certificates information from the EDP civil registry data. We compare the birth-certificate information (e.g., gender, parents' age at birth, single parenthood, parents' occupation at birth) for all EDP individuals born in 1972-1981 in metropolitan France and for our sample of children. Note that the resulting statistics are subject to the imperfections of birth-certificate data, notably regarding non-random missing information for fathers. Table C.1 displays the statistics for both samples. Overall, our sample of children is very similar to the unrestricted EDP sample, except for a higher probability of being in the fiscal data (91% vs. 100%, by construction) and a lower likelihood of having a father who is a farmer. The household income distributions are very similar.

Characteristic	Population	Sample	Diff.
Females	49.14%	49.77%	0.63
<i>Parent demographics</i>			
Mother age at birth	26	25.89	-0.11
Father age at birth	28.91	28.65	-0.26
Mother born French	90.07%	91.92%	1.85
Father born French	88.12%	90.15%	2.03
Single mothers	4.98%	4.42%	-0.56
Missing parents info.	2.24%	1.75%	-0.49
<i>Father 1-digit occupation at child birth</i>			
Missing father info.	9.4%	8.25%	-1.15
1. Farmers	3.41%	0.64%	-2.77
2. Craftsmen, salespeople, and heads of businesses	3.95%	3.96%	0.01
3. Managerial and professional occupations	7.14%	5.98%	-1.16
4. Intermediate professions	13.58%	14.63%	1.05
5. Employees	14.58%	16%	1.42
6. Blue collar workers	46.46%	49.4%	2.94
7. Retirees	0.03%	0.02%	-0.01
8. Other with no professional activity	1.45%	1.11%	-0.34
<i>Mother 1-digit occupation at child birth</i>			
Missing mother info.	5.34%	4.69%	-0.65
1. Farmers	0.83%	0.11%	-0.72
2. Craftsmen, salespeople, and heads of businesses	0.91%	0.85%	-0.06
3. Managerial and professional occupations	2.08%	1.62%	-0.46
4. Intermediate professions	8.92%	9.19%	0.27
5. Employees	26.19%	28.33%	2.14
6. Blue collar workers	11.2%	12%	0.8
7. Retirees	0.02%	0.02%	0
8. Other with no professional activity	44.51%	43.2%	-1.31
<i>All Employee Panel (AEP) information in adulthood, 1968-2015, age 35-45</i>			
Observed in AEP	72.83%	78.67%	5.84
Mean number of obs. in AEP	2.9	3.15	0.25
Q1 individual wage (AEP)	12,671	13,179	508
Mean individual wage (AEP)	21,538	21,666	128
Med. individual wage (AEP)	19,528	19,726	198
Q3 individual wage (AEP)	26,623	26,723	100
<i>Tax information in adulthood, 2010-2016, age 35-45</i>			
Observed in tax data	90.92%	100%	9.08
Mean number of obs. in tax data	4.23	4.65	0.42
Q1 household income (tax)	27,339	27,696	357
Mean household income (tax)	46,858	46,598	-260
Med. household income (tax)	41,220	41,418	198
Q3 household income (tax)	56,630	56,481	-149
N	83,009	64,571	

Notes: Comparison of birth-certificate information on the full EDP sample vs. the study sample. See Section 3.2 for details on construction of the study sample.

Table C.1: Average Characteristics of Overall Population vs. Sample

C.2. *Alternative First-Stage Estimation*

The parent income predictions we use to palliate French data limitations are central to our analysis. It is of primary importance that the first stage of the two-step strategy we rely on is reliable. We make sure that this first stage does not spuriously drive the results in one way or another by evaluating its sensitivity to varying the set of instruments and to relaxing parametric assumptions.

Set of First-Stage Predictors

The most important dimension to consider is the set of variables included in the first stage, notably because it has been shown that inadequate instruments could yield inconsistent estimates (Jerrim et al., 2016). Appendix Figure C.2 documents the sensitivity of IGE, RRC and transition matrix estimates to the set of predictors used in the first-stage estimation. We estimate them when adding each of the following predictors sequentially (all measured in 1990): education (8 categories), 2-digit occupation (42 cat.), a group of demographic characteristics (age, French nationality dummy, country of birth (6 cat.), and household structure (6 cat.)) and a group of municipality-level characteristics (unemployment rate, share of single mothers, share of foreigners, population, and population density). Since relying on a single variable with less than 100 categories induces some income values to span over several percentiles, parents with a given predicted income are attributed the average rank of individuals earning that level of income. Lastly, we also report the adjusted R^2 , computed as the average from 5-fold cross-validation.

We find that the IGE is 0.68 when using only education as the first-stage predictor, consistent with a point already made in the literature that using only education as a predictor is likely to yield inflated estimates of the IGE. Once 2-digit occupation is included in the first-stage, adding other demographic or municipality-level characteristics has no effect on the estimates. Indeed, as can be seen from the R^2 , most of the predictive power actually comes from the 2-digit occupation variable. The RRC appears remarkably unchanged by the set of first-stage predictors used, at 0.28 with only education and 0.30 with all variables. This appears once more to be a strength of the RRC in the TSTSLs context.

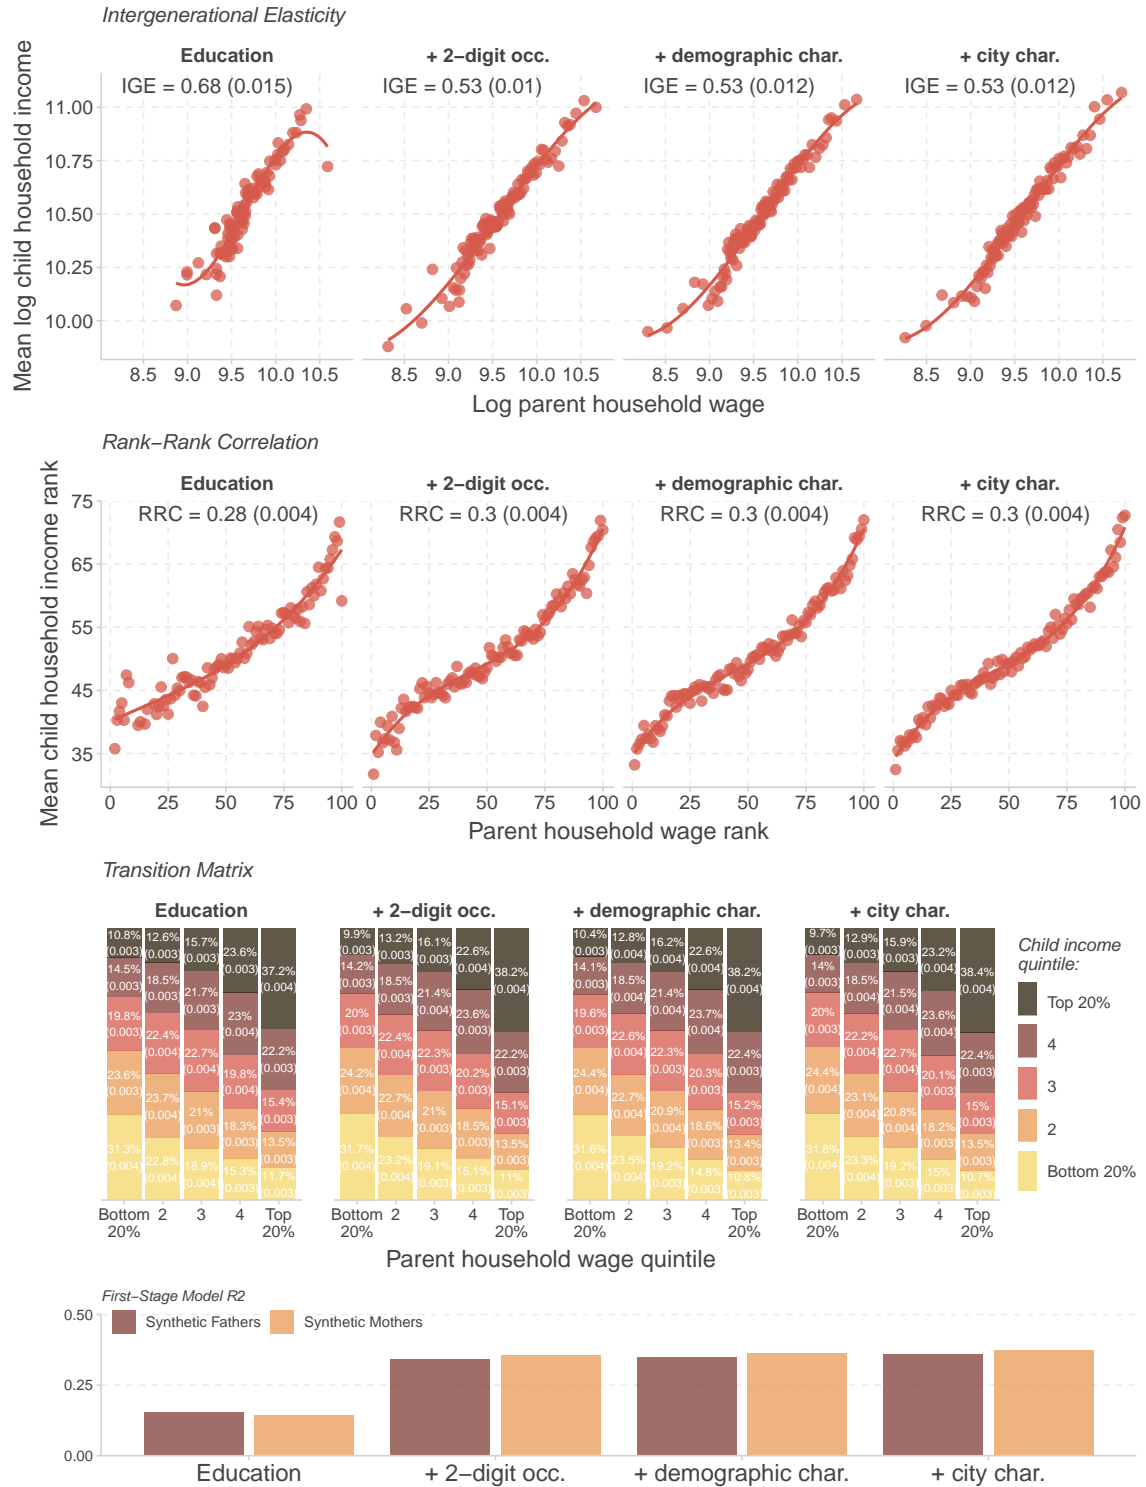


Figure C.2. Robustness of Baseline Estimates to Different First-Stage Predictors

Notes: This figure assesses the robustness of our baseline IGE, RRC and transition matrix estimates to variations in the set of first-stage predictors. Parent income is predicted separately for fathers and mothers using a set a of instruments that vary along the x-axis. We report the corresponding CEFs, along with the point estimates and the bootstrapped standard error in parenthesis. The bottom panel of the figure reports separately for synthetic fathers and mothers the R^2 associated with each first stage. See Figure 1.3's notes for details on data, sample and income definitions.

Flexible Models

We make use of semi- and non-parametric models to elicit potential misspecifications in the first stage. The baseline specification of the first stage is of the form $y = \beta X + \varepsilon$, where y is the log of parent lifetime income and X is a set of k predictors. OLS would not account for interactions between predictors nor for non-linearities in the relationship between X and y unless they are explicitly modeled. Fully non-parametric methods of the form $y = m(X) + \varepsilon$ would capture both interactions and non-linearities that may help reduce the out-of-sample MSE. Obtaining a lower MSE and significantly different second-stage estimates with non-parametric models than with OLS would suggest that non-modeled non-linearities, interactions, or both, influence the resulting intergenerational mobility estimates.

We implement this test using three machine learning methods: (i) a generalized additive model (GAM) of the form $y = m_1(x_1) + m_2(x_2) + \dots + m_k(x_k) + \varepsilon$ which accounts for non-linearities but not for interactions unless explicitly specified, (ii) a gradient boosted regression tree, that is a high-dimensional combination of sequentially grown regression trees, and (iii) the ensemble method, which consists in taking the average of the predictions from each model weighted in a way that minimizes the out-of-sample MSE.

Appendix Figure C.3 compares the intergenerational mobility estimates and out-of-sample MSE resulting from these three methods using our baseline child and parent income definitions. We do not observe significant differences in MSE between the different prediction methods. The resulting mobility estimates are virtually the same for OLS, GAM and the ensemble method, and slightly smaller for boosted trees. This suggests that conditional on the set of predictors we use, using more flexible estimation methods does not lead to better income predictions and different estimates than using an additive OLS specification.




Figure C.3. Robustness to Machine Learning Prediction

Notes: This figure assesses the robustness of our baseline intergenerational income mobility estimates to increasingly flexible first-stage prediction models. Each bar represents the magnitude of the estimate of the corresponding color estimated using the first-stage model indicated on the x-axis. The first set of estimates are the baseline estimates obtained using OLS. The three other sets are obtained using increasingly flexible models: generalized additive models (GAM), gradient boosted regression trees, and the ensemble method. The connected dots represent the average out-of-sample MSEs of the associated prediction models, estimated using 5-fold cross-validation. See Figure 1.3’s notes for details on data, sample and income definitions.

C.3. Lifecycle and Attenuation Bias

Child Lifecycle Bias - Constant Sample of Children

To overcome the issue related to changes in Figure 1.6’s underlying sample of children, we reproduce the individual wage estimates using the All Employee Panel keeping the sample of children constant. To do so we restrict to children born in 1972 and 1974⁴⁵ for whom wages are observed every year between 25 and 43 years old and 25 and 41 years old respectively. Appendix Figure C.4 displays the results. Since the sample is kept constant throughout, the coefficients can be compared to one another and the change in magnitude can only be driven by the age at which child income is measured rather than sample composition. As in Figure 1.6, we find that measuring child income prior to the mid-thirties seriously underestimates the IGE (panel A) and RRC (panel B), and overestimates (underestimates) bottom or top mobility (persistence) (panel C).

⁴⁵We cannot include the 1973 cohort as the All Employee Panel income data are only available for individuals born an even year before 2001. This choice of cohorts is done to be able to measure their incomes after they are 40 years old.

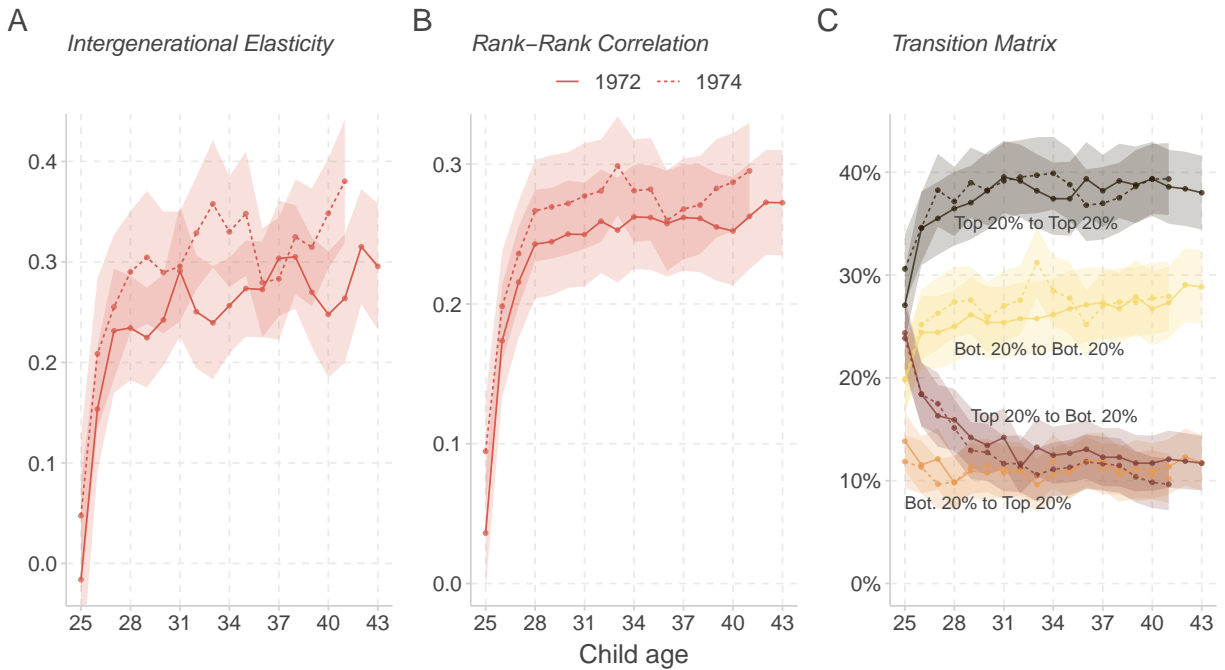



Figure C.4. Child Lifecycle Bias - 1972 and 1974 Cohorts (Constant Sample)


Notes: This figure assesses the robustness of our baseline intergenerational income mobility estimates presented in Figures 1.3 and 1.4 to changes in the age at which child income is measured, for children born in 1972 (solid line) and 1974 (dashed line). For both birth cohorts the sample is kept constant, that is only children with wages observed in the All Employee Panel at each age between 25 and 43 years old are retained. Shaded areas represent the 95% bootstrapped confidence interval. See Sections 3 and 4.4 for details on data, sample and income definitions.

Child and Parent Lifecycle Bias Jointly

Child and parent lifecycle bias are typically assessed independently, as we do in the main body of the article. Yet they influence one another and it is instructive to estimate our measures of intergenerational persistence for each possible combination of synthetic parent and child age. Appendix Figure C.5 shows such estimates when child income is measured between ages 30 and 44, and synthetic parent income between ages 28 and 60.



(a) Intergenerational Elasticity



(b) Rank-Rank Correlation

Figure C.5. Child and Parent Lifecycle Bias

Notes: This figure assesses the robustness of our baseline intergenerational income mobility estimates presented in Figure 1.3 to changes in the age at which child and synthetic parent incomes are measured. The sample of children and synthetic parents varies across ages. See Sections Figure 1.3's notes for details on data, sample and income definitions.

Parent Attenuation Bias

Figure C.6 plots estimates of our persistence measures varying the number of synthetic parent income observations used in the first-stage regression from 1 to 11. To control for the potential effect of lifecycle bias we center the age at which synthetic parent income is measured at 40 years old. In other words, one income observation corresponds to income at age 40, two income observations corresponds to average income at ages 39 and 41, three income observations to average income between 39 and 41, and so on. Therefore, 11 income observations corresponds to the average between 35 and 45 years old. The sample of synthetic parents over which income is predicted varies for each estimate depending on how many synthetic parents had incomes observed each year in the required age range. We report results both for parent household wage and father wage.

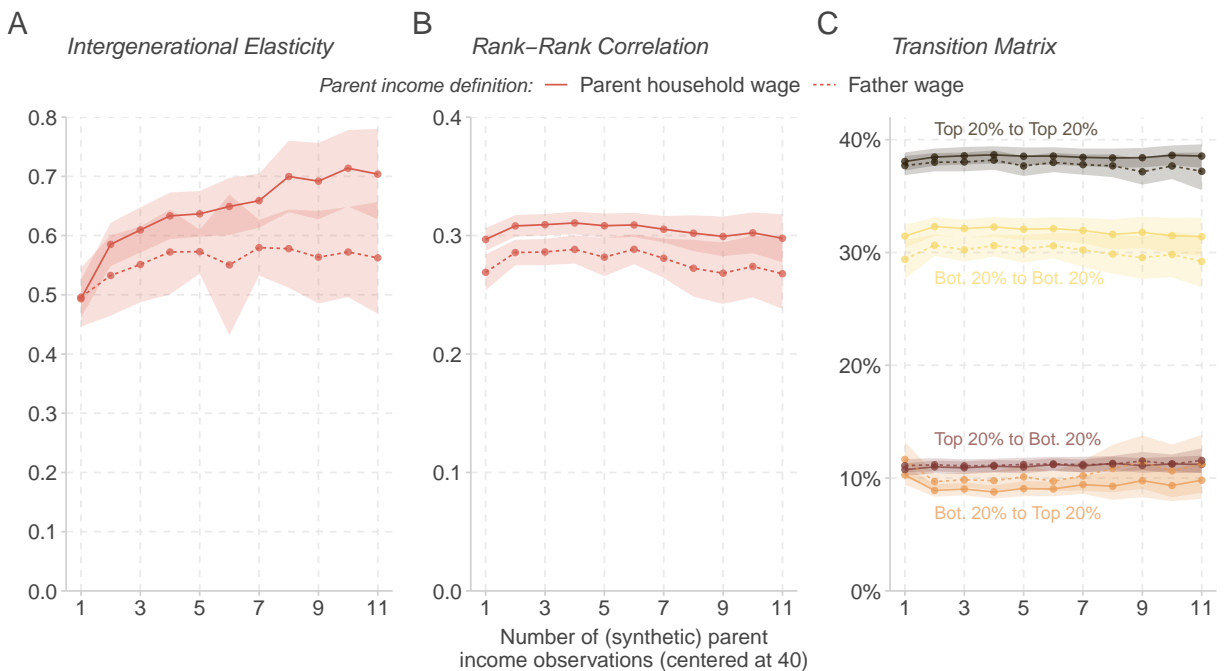


Figure C.6. Attenuation Bias

Notes: This figure assesses the robustness of our baseline intergenerational income mobility estimates to the number of income observations used to predict parent income. While varying the number of parent income observations, we center the age range at 40 to control for lifecycle bias. Shaded areas represent the 95% bootstrapped confidence interval. See Figure 1.3's notes for details on data, sample and income definitions.

These results suggest that attenuation bias might affect our baseline IGE (panel A) but not our other estimates of intergenerational mobility. Indeed when defining parent income at the household level, the IGE increases from 0.5 when using only one income observation to 0.7 when averaging over 11 income observations (i.e., between 35 and 45). It is important to highlight that almost all of this change is driven by how mothers' incomes are predicted. Indeed when looking at the father-child IGE, the estimate does not increase so markedly and stabilizes around 2 or 3 income observations, consistent with the idea that the two-stage procedure employed drastically shrinks the transitory component of annual income, and in large contrast with what is typically

found when parent income is actually observed (Mazumder, 2005). Indeed, since we are already predicting parent income based on observable characteristics, and thus in a sense reducing year-to-year income volatility, averaging over more years does not affect the estimate much.

How one interprets the results based on parent household wage depends on one's prior as to how to best predict mothers' incomes. Our view is that predicting mothers' incomes only on the subsample of synthetic mothers with observed wages in all years between 35 and 45 years old biases the underlying sample considering the uneven labor force participation of women at the time. We believe our choice of restricting our sample of synthetic parents to those with at least two income observations between ages 35 and 45 is reasonable.

Constant Sample. We check whether the lack of change in intergenerational mobility measures with the number of synthetic parent income observations observed in Figure C.6 could be due to the fact that the sample of synthetic parents varies throughout. We replicate those estimates restricting the sample of synthetic parents to those with all 11 income observations between 35 and 45 years old and estimating the intergenerational mobility measures by varying the number of income observations averaged in the first-stage regression (centered around 40 years old again). To do so, we impute wages in 1981, 1983 and 1990, for which the data are not available,⁴⁶ using the average wage between the previous and subsequent year only if both wages are observed. This enables us to have a consistent sample and increase the number of synthetic parents on which the predictions can be done.

Appendix Figure C.7 displays the results from this sensitivity analysis. The increase in the parent household wage IGE is much less marked, increasing from 0.636 when using one income observation to 0.704 when using all 11 observations (panel A). Our interpretation of this relatively modest increase is that averaging over at least 2 income observations as we do for our baseline estimate should suffice to not suffer from attenuation bias. Note that what matters in this figures is not how different the estimates are from our baseline estimate but rather the extent to which they vary with the number of synthetic parent income observations used. Indeed, the difference between our baseline IGE estimate and the estimates obtained are driven by the fact that the sample of synthetic parents for whom we observe all incomes between 35 and 45 years old is a highly non-representative sample, especially when it comes to mothers. In fact, we do not find any attenuation bias when restricting our analysis to fathers, suggesting all the variation in the IGE can be accounted for by changes in mothers' incomes predictions. As with the varying synthetic parent sample estimates, rank-based intergenerational mobility measures are significantly less sensitive to averaging over more income years, and the estimates found are very close to our baseline ones (panels B and C).

⁴⁶As explained in Appendix A, the 1982 and 1990 population censuses generated an extra workload which prevented INSEE from compiling the All Employee Panel data for these years.

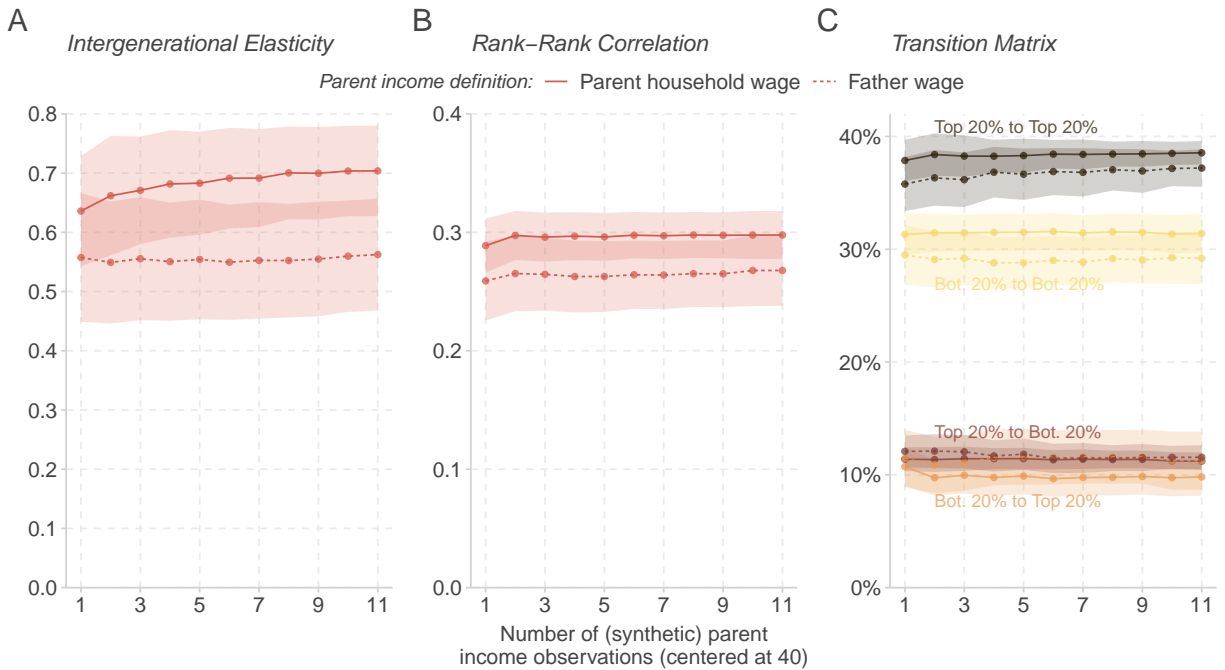


Figure C.7. Parent Attenuation Bias (Constant Sample of Synthetic Parents)

Notes: This figure assesses the robustness of our baseline intergenerational income mobility estimates to the number of income observations used to predict parent income, keeping the sample of synthetic parents constant. The sample of synthetic parents is thus restricted to those with all 11 income observations between 35 and 45 years old. While varying the number of parent income observations, we center the age range at 40 to control for lifecycle bias. Shaded areas represent the 95% bootstrapped confidence interval. See Figure 1.3's notes for details on data, sample and income definitions.

Child Attenuation "Bias"

Appendix Figure C.8 plots estimates of our persistence measures varying the number of child income observations from 1 to 7 between 35 and 45 years old, keeping the sample of children constant⁴⁷ (i.e. keeping only children with 7 household income observations). Due to this restriction only cohorts born between 1972 and 1975 are kept. Without this restriction, the value reported for 1 income observation would correspond to our baseline estimate. In the same way as for parents, we control for lifecycle bias by centering the year in which child income is measured to 2013. In other words, one child income observation corresponds to income measured in 2013, two income observations corresponds to the average between 2012 incomes and 2014 incomes, three to average income between 2012 and 2014, etc. The results suggest that estimates are largely unaffected by increasing the number of child income observations.

⁴⁷The sample varies ever so slightly for the IGE due to the number of negative or 0 incomes changing between years.




Figure C.8. Sensitivity to Number of Child Income Observations (Constant Sample)

Notes: This figure presents estimates of our persistence measures varying the number of child income observations from 1 to 7 between 35 and 45 years old, keeping the sample of children constant, i.e. keeping only children with 7 household income observations. (The sample varies ever so slightly for the IGE due to the number of negative or 0 incomes varying between years.) Due to this restriction only cohorts born between 1972 and 1975 are kept. Without this restriction, the value reported for 1 income observation would equal our baseline estimate. We control for lifecycle bias by centering the year in which child income is measured to 2013. In other words, one child income observation corresponds to income measured in 2013, two income observations corresponds to the average of 2012 and 2014, three to average income between 2012 and 2014, etc. Shaded areas represent the 95% bootstrapped confidence interval. See Figure 1.3’s notes for details on data, sample and income definitions.

C.4. Sensitivity to Income Distribution Tails.

Our baseline estimates may be sensitive to two main sample selection choices when it comes to the income distributions of parent and children: (i) how children with negative or zero incomes are treated; and (ii) how the top and bottom tails of both the parent and child income distributions are dealt with.

Treatment of Zeros

The first issue is particularly salient for the estimation of the intergenerational income elasticity due to the impossibility of taking the log of zero.⁴⁸ Many researchers simply discard such observations since they are likely not representative of lifetime income. Though this may potentially be the case if only short income time spans are available, we nonetheless evaluate how our baseline

⁴⁸Various methods have been proposed to overcome this issue. Bellégo et al. (2021) describe such methods and propose a novel solution that can be applied to a variety of cases.

estimates of both the IGE and the RRC when replacing negative or zero child income values by 1 or 1,000 euros.

Appendix Figure C.9 shows estimates for the IGE and RRC when replacing income of children reporting negative or zero incomes by 1 euro or 1,000 euros, for different child income definitions. For our primary child income definition, household income, the estimates do not change due to there being very few children with negative or zero household income. However, for child income defined at the individual-level, for which the share of negative or zero incomes is more important, the IGE becomes highly sensitive to the recoding of such observations while the RRC remains unchanged. For example, for individual child income, the IGE is 0.46 when zeros are dropped and 0.82 when they are recoded to 1 and 0.55 when recoded to 1,000. The RRC is entirely insensitive to such recoding as ranks are not altered by it.




Figure C.9. Sensitivity to Different Zero Child Income Replacement Values

Notes: This figure assesses the robustness of our baseline IGE and RRC estimates to replacing incomes of children reporting negative or zero incomes by 1 euro or 1,000 euros, for different child income definitions. Vertical lines represent the 95% bootstrapped confidence intervals. See Section 3 for details on data, sample and income definitions.

Top and Bottom Trimming

The second issue relates to the treatment of top and bottom earners in both the parent and child income distributions. For the parent income distribution the choice can both be made in the prediction stage and in the second stage. Specifically, we assess how the IGE and RRC vary when trimming the top and/or bottom 1% to 5% and 10%. Appendix Figure C.10 displays the results of


this sensitivity check. There are three main takeaways.

First, the IGE is significantly more sensitive to small changes in parent or child income distributions while the RRC remains relatively stable. For example, removing the top and bottom 1% of child incomes decreases the IGE from 0.527 to 0.418 while the RRC only decreases from 0.303 to 0.286. It does not seem desirable that a measure of intergenerational mobility should be so sensitive to excluding just 2% of children. Mathematically it can be linked to changes in the dispersion of the distribution of child incomes but conceptually it seems difficult to defend such responsiveness to minor sample changes.

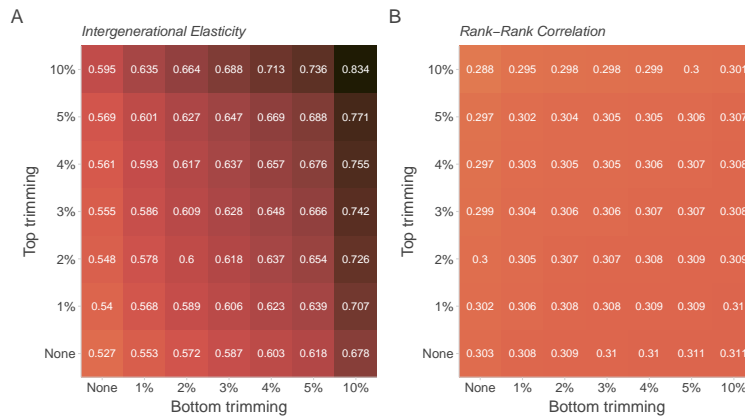
Second, the IGE is quite strongly influenced by minor trimming in the first-stage prediction sample. For example, excluding the bottom and top 2% of synthetic parent incomes leads to an IGE of 0.6. Such exclusions are not uncommon in the literature though their relevance is unclear.⁴⁹ Meanwhile the RRC is once more remarkably untouched by first-stage parent income exclusions. In fact excluding the bottom and top 10% of synthetic parent incomes decreases the RRC to 0.301 (from 0.303). This appears to be an additional benefit of estimating the RRC when using with the TSTSLs method. Note that trimming the first-stage prediction sample does lead to increased out-of-sample MSE, as shown in Appendix Figure C.11.

Third, for second-stage parent income trimming, the effects are relatively mild for both intergenerational mobility measures. This is very likely a consequence of the two-stage procedure which reduces the variance in parent incomes.

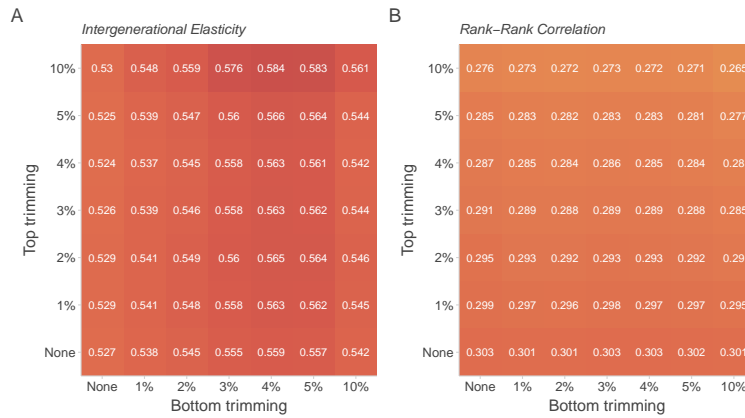
⁴⁹For example, Barbieri et al. (2020) exclude the top and bottom 1% of their sons and synthetic fathers' incomes.



(a) Child Income Trimming



(b) First-Stage Synthetic Parent Income Trimming



(c) Second-Stage Parent Income Trimming

Figure C.10. Sensitivity to Child and Parent Income Distributions Trimming

Notes: This figure assesses the robustness of our baseline intergenerational income mobility estimates presented to trimming the tails of the parent and child income distributions. Each cell displays the value of the corresponding intergenerational mobility measure obtained after trimming the income distribution of the corresponding sample by the fraction indicated on the x-axis at the bottom and by that indicated on the y-axis at the top. See Figure 1.3's notes for details on data, sample and income definitions.




Figure C.11. Out-of-sample MSE as a Function of Top and Bottom Trimming

Notes: This figure plots the out-of-sample MSE as a function of trimming various shares of the tails of synthetic parents' income distribution. Our-of-sample MSEs correspond to the average MSE obtained from 5-fold cross-validation. See Sections 3.1 and 3.2 for details on the exact model being estimates and sample construction.

C.5. Transition Probabilities at the Top

To analyze persistence at the top of the parent income distribution, we estimate transition matrices for the top 10%, top 5% and top 2% of parent incomes and compare our results with those from the United States.⁵⁰ We estimate the likelihood of remaining in the top 10% to be about 28% in France close to the United States figure of 26%. This statistic is almost 3 times larger than would be observed in a world where child income is unrelated to parent income (i.e., 10%). This persistence at the top gets stronger as we zoom into the top 5% (22% remaining in top 5%) and top 2% (14% remaining in top 2%). The ratio of observed persistence to counterfactual world with no link between incomes increases with parent income rank in the distribution. This suggests that mechanisms of intergenerational persistence at the top of the parent income distribution might differ from those at play for the rest of the distribution.

⁵⁰We use the detailed percentile-by-percentile estimates provided in the online appendix of Chetty et al. (2014).

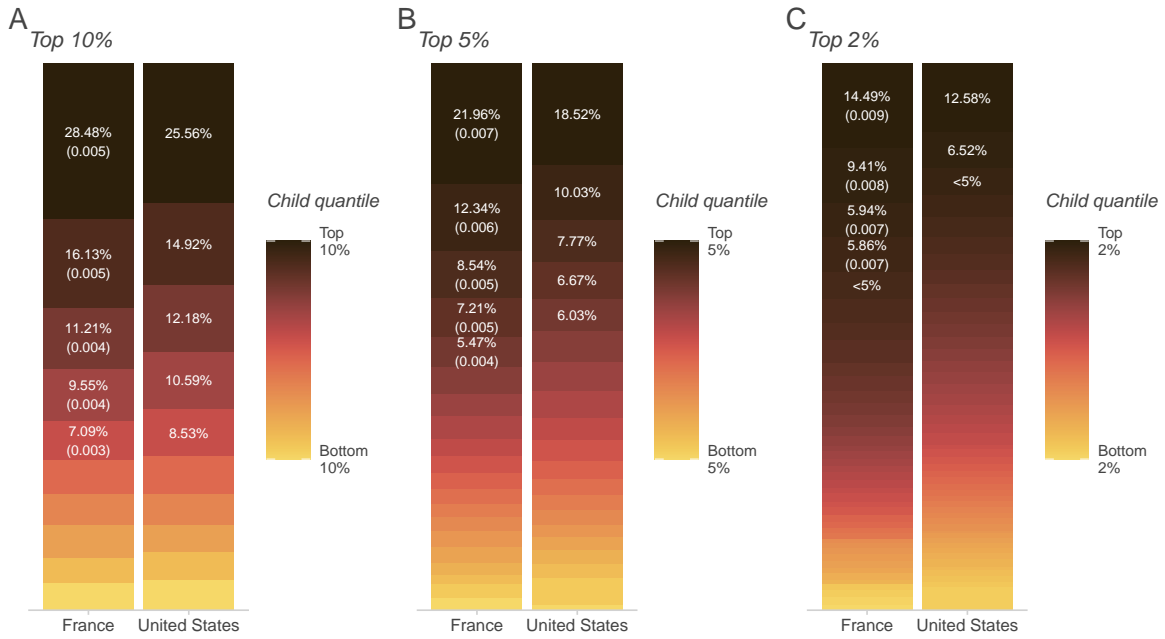


Figure C.12. Top Parent Income Quantiles Transition Matrices in France and United States

Notes: This figure presents intergenerational transition matrix estimates for children coming from families in the top 10% (panel A), top 5% (panel B) and top 2% (panel C) of the parent income distribution, with bootstrapped standard errors in parentheses. We compare the transition probabilities we obtain for France with those computed by [Chetty et al. \(2014\)](#) for the United States. Each cell corresponds to the percentage of children in a given income quantile who have parents in a given parent income quantile. See Section 3 for details on data, sample and income definitions.

D. Correlation with Local Characteristics

D.1. Definitions and Data Sources

Appendix Table D.1 displays the variables used in the correlational analysis presented in Section 6 (subsection *Correlation with Local Characteristics*). We measure these variables as close to 1990 as possible so as to reflect the environment individuals grew up in.

Variable	Definition	Source
Demographic		
Density	Log number of inhabitants per square meter	1990 BDCOM ¹
% Foreigner	Share without French nationality	1990 Census
% Single mothers	Share of single mothers in the adult population (≥ 18)	1990 Census
Economic		
Mean wage	Log average wage	1996 DADS Panel
% Unemployed	Unemployment rate	1990 Census
Inequality		
Gini index	Gini index of wage inequality	1996 DADS Panel
Theil index	Theil index of spatial wage segregation	1996 DADS Panel
Share top 1%	Share of total wage accrued by the top 1% of wage earners	1996 DADS Panel
Education		
# HEI	Number of higher education institutions	2007 BPE ²
Distance to HEI	Average distance to the closest public higher education institution	2007 BPE ²
% HS graduates	Share of high-school graduates in adult population (≥ 18)	1990 Census
Social capital		
Cultural amenities	Number of cinemas and museums per capita	2007 BPE ² , Min. de la Culture
Crime	Number of offenses and crimes per capita	Min. de l'Intérieur
% Voters	Participation rate to the first round of the 1995 presidential election	Min. de l'Intérieur

Notes:

¹ Base de données communales du recensement de la population (BDCOM) - 1990, INSEE (producteur), ADISP (diffuseur) - doi:10.13144/lil-0363

² Base permanente des équipements (BPE) - 2007, INSEE (producteur), PROGEDO-ADISP (diffuseur) - doi:10.13144/lil-0423

Table D.1: Definitions and Sources of Department Characteristics

D.2. Simple Regression Analysis

We start by regressing department-level intergenerational mobility estimates on each of these variables in separate regressions. Both the department intergenerational mobility estimates and the characteristics are standardized, implying that the coefficients can be interpreted as correlations. Results are presented in Appendix Tables D.2 to D.4 and summarized in Figure 1.9. Note that for the IGE and RRC, a positive coefficient implies the characteristic is positively correlated with

intergenerational *persistence* (i.e., negatively correlated with intergenerational *mobility*), while for absolute upward mobility a positive coefficient implies the characteristic is positively correlated with higher incomes for children born to low-income families.

Appendix Figure D.1 provides a potential explanation for the results of the correlational analysis by documenting the correlation between all department characteristics. The 14 variables considered are for the most part quite strongly correlated with one another, both within and between variable groups. For instance, the Gini index is highly correlated with other inequality measures, but also with population density and the share of high school graduates, two variables whose relationship with absolute upward mobility is positive.




Figure D.1. Correlation Between Department Characteristics

Notes: This figure presents the correlation coefficient between all department characteristics considered, defined as follows. See Appendix Table D.1 for definitions and sources of the department characteristics.

D.3. Lasso Analysis

Considering the strong correlation across department characteristics, we estimate lasso regressions in order to identify the characteristics that are most strongly associated with intergenerational mobility. The result of this analysis is presented in Appendix Figure D.2.

The lasso analysis does not change the picture much. For the IGE, only the unemployment rate is picked up, as was the case in the univariate setting. For the RRC, the lasso maintains some demographic characteristics (% of single mothers and % foreigners), the unemployment rate, all three education variables, and two measures of social capital (cultural amenities and % voters). Again, these results are largely in line with what was observed in the univariate regressions. Lastly, for absolute upward mobility roughly the same characteristics that were significant in the simple regression analysis are kept except importantly for the Gini index.

Though the relationships we document between intergenerational mobility and department characteristics are overall pretty intuitive, these descriptive relationships cannot distinguish a potential causal effect of place from sorting. We leave this causal assessment to future studies.




Figure D.2. Department Characteristics Kept by Lasso

Notes: This figure presents the department characteristics kept by the lasso regression. See Appendix Table D.1 for definitions and sources of the department characteristics.

	Dependent variable: Intergenerational Elasticity													
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)	(13)	(14)
Density	-0.022 (0.110)													
% Single mothers		0.016 (0.110)												
% Foreigners			-0.026 (0.110)											
Unemployment rate				0.306*** (0.104)										
Mean wage					-0.157 (0.108)									
Distance to HEI						-0.100 (0.109)								
% HS graduates							-0.114 (0.109)							
# HEI								-0.022 (0.110)						
Theil index									-0.024 (0.110)					
Share top 1%										-0.092 (0.109)				
Gini index											0.007 (0.110)			
Cultural amenities												-0.090 (0.109)		
Crime													0.086 (0.109)	
% Voters														0.042 (0.110)
Intercept	4.124*** (0.135)	4.008*** (0.893)	4.183*** (0.213)	2.512*** (0.565)	18.162* (9.654)	4.374*** (0.277)	4.554*** (0.412)	4.218*** (0.406)	4.397*** (1.169)	4.954*** (0.977)	4.000* (2.287)	4.386*** (0.318)	3.911*** (0.312)	2.943 (3.152)
Observations	85	85	85	85	85	85	85	85	85	85	85	85	85	85
R ²	0.0005	0.0003	0.001	0.094	0.025	0.010	0.013	0.0005	0.001	0.008	0.00005	0.008	0.007	0.002

Notes: All variables are standardized such that the regression coefficient corresponds to the correlation. See Appendix Table D.1 for variable definitions and data sources. Standard errors in parentheses. *p<0.1; **p<0.05; ***p<0.01.

Table D.2: Correlation Between Intergenerational Elasticity and Department Characteristics

	Dependent variable: Rank-Rank Correlation													
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)	(13)	(14)
Density	-0.083 (0.109)													
% Single mothers		-0.168 (0.108)												
% Foreigners			-0.255** (0.106)											
Unemployment rate				0.181* (0.108)										
Mean wage					-0.131 (0.109)									
Distance to HEI						-0.078 (0.109)								
% HS graduates							-0.131 (0.109)							
# HEI								0.099 (0.109)						
Theil index									0.055 (0.110)					
Share top 1%										-0.105 (0.109)				
Gini index											0.008 (0.110)			
Cultural amenities												-0.141 (0.109)		
Crime													-0.049 (0.110)	
% Voters														0.239** (0.107)
Intercept	5.197*** (0.135)	6.617*** (0.881)	5.683*** (1.206)	4.294*** (0.583)	16.940* (9.692)	5.440*** (0.278)	5.736*** (0.411)	4.906*** (0.404)	4.675*** (1.167)	6.186*** (0.976)	5.097** (2.287)	5.642*** (0.316)	5.389*** (0.312)	-1.612 (3.064)
Observations	85	85	85	85	85	85	85	85	85	85	85	85	85	85
R ²	0.007	0.028	0.065	0.033	0.017	0.006	0.017	0.010	0.003	0.011	0.0001	0.020	0.002	0.057

Notes: All variables are standardized such that the regression coefficient corresponds to the correlation. See Appendix Table D.1 for variable definitions and data sources. Standard errors in parentheses. *p<0.1; **p<0.05; ***p<0.01.

Table D.3: Correlation Between Rank-Rank Correlation and Department Characteristics

	Dependent variable: Absolute Upward Mobility													
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)	(13)	(14)
Density	0.491*** (0.096)													
% Single mothers		0.218** (0.107)												
% Foreigners			0.478*** (0.096)											
Unemployment rate				-0.563*** (0.091)										
Mean wage					0.591*** (0.089)									
Distance to HEI						-0.293*** (0.105)								
% HS graduates							0.530*** (0.093)							
# HEI								0.274** (0.106)						
Theil index									0.337*** (0.103)					
Share top 1%										0.406*** (0.100)				
Gini index											0.288*** (0.105)			
Cultural amenities												0.015 (0.110)		
Crime													0.221** (0.107)	
% Voters														-0.405*** (0.100)
Intercept	13.426*** (0.118)	11.305*** (0.872)	12.268*** (0.187)	16.055*** (0.490)	-39.580*** (7.885)	13.751*** (0.266)	11.138*** (0.352)	12.092*** (0.390)	9.488*** (1.100)	9.459*** (0.896)	7.068*** (2.190)	13.026*** (0.319)	12.476*** (0.305)	24.709*** (2.884)
Observations	85	85	85	85	85	85	85	85	85	85	85	85	85	85
R ²	0.241	0.048	0.229	0.317	0.349	0.086	0.280	0.075	0.114	0.165	0.083	0.0002	0.049	0.164

Notes: All variables are standardized such that the regression coefficient corresponds to the correlation. See Appendix Table D.1 for variable definitions and data sources. Standard errors in parentheses. *p<0.1, **p<0.05, ***p<0.01.

Table D.4: Correlation Between Absolute Upward Mobility and Department Characteristics

E. Additional Figures




Figure E.1. Family Position in 1990 Census by Child Birth Cohort

Notes: This figure presents the family position of EDP individuals in the 1990 census by birth cohort. The sample is restricted to EDP individuals born in metropolitan France.




Figure E.2. Baseline IGE Estimates for Different Child and Parent Income Definitions

Notes: This figure presents our baseline intergenerational income elasticity estimates for various parent and child income definitions. Each bar represents the coefficient of an OLS regression of child income on parent income, for the entire sample (All) and for sons and daughters separately. Error bars represent the 95% bootstrapped confidence intervals. See Section 3 for details on data, sample and income definitions.




Figure E.3. Baseline RRC Estimates for Different Child and Parent Income Definitions

Notes: This figure presents our baseline intergenerational rank-rank correlation estimates for various parent and child income definitions. Each bar represents the coefficient of an OLS regression of child income rank on parent income rank, for the entire sample (*All*) and for sons and daughters separately. Error bars represent the 95% bootstrapped confidence intervals. See Section 3 for details on data, sample and income definitions.




Figure E.4. Baseline Quintile Transition Matrix for Different Child Income Definitions

Notes: This figure presents our baseline intergenerational transition matrix estimates for various child income definitions, with bootstrapped standard errors in parentheses. Each cell corresponds to the percentage of children in a given income quintile among children who have parents in a given parent income quintile. See Section 3 for details on data, sample and income definitions.




Figure E.5. Rank-Rank Correlation and Upward Mobility in International Comparison

Notes: This figure represents the international comparisons in rank-rank correlation and transition matrix cells presented in Tables 1.1 and 1.2.




Figure E.6. Higher Education Graduation by Quintile Transition Matrix Cell

Notes: This figure presents the percentage of children graduating from higher education in each cell of the quintile transition matrix. Each cell corresponds to the percentage of children in a given income quintile coming from a family in a given parent income quintile who have a higher education diploma. See Sections 3 and 4.4 for details on data, sample and income definitions. In this figure parent income ranks are computed without parent education in the set of first-stage predictors to avoid capturing the effect of parent education independent from that of parent income.



Figure E.7. French Departments

Notes: This figure represents the 96 metropolitan French departments. The borders of these departments have not changed over the study period. For convenience, we treat Corsica (*Haute Corse* and *Corse du Sude*) as a single department.




Figure E.8. Illustration of Absolute Upward Mobility for the *Nord* Department

Notes: This figure presents a non-parametric binned scatter plot of the relationship between child income rank and parent income rank for individuals who grew up in the *Nord* department. The dashed line shows the expected income rank, here 38.7 (bootstrapped standard error = 0.54), for children whose parents locate at the 25th percentile. The orange line is a linear regression fit through the conditional expectation. See Figure 1.3's notes for details on data, sample and income definitions.




Figure E.9. Department-Level Log-Log Relationships

Notes: This figure presents the non-parametric binned scatter plot of the relationship between child log income and parent log income separately for each childhood department. The childhood department is that observed in 1990, i.e., when individuals were between 9 and 18 years old. See Figure 1.3's notes for details on data, sample and income definitions.




Figure E.10. Department-Level Rank-Rank Relationships

Notes: This figure presents the non-parametric binned scatter plot of the relationship between child income rank and parent income rank separately for each childhood department. The childhood department is that observed in 1990, i.e., when individuals were between 9 and 18 years old. See Figure 1.3’s notes for details on data, sample and income definitions.




Figure E.11. Department-Level Intergenerational Elasticities

Notes: This figure presents the intergenerational elasticity in household income and its 95% bootstrapped confidence interval, estimated separately for each childhood department with more than 200 observations. The childhood department is that observed in 1990, i.e., when individuals were between 9 and 18 years old. The dashed black line represents the national estimate. See Figure 1.3's notes for details on data, sample and income definitions.




Figure E.12. Department-Level Rank-Rank Correlations

Notes: This figure presents the rank-rank correlation in household income and its 95% bootstrapped confidence interval, estimated separately for each childhood department with more than 200 observations. The childhood department is that observed in 1990, i.e., when individuals were between 9 and 18 years old. The dashed black line represents the national estimate. See Figure 1.3's notes for details on data, sample and income definitions.




Figure E.13. Department-Level Absolute Upward Mobility

Notes: This figure presents the absolute upward mobility in household income ranks and its 95% bootstrapped confidence interval, estimated separately for each childhood department with more than 200 observations. The childhood department is that observed in 1990, i.e., when individuals were between 9 and 18 years old. The dashed black line represents the national estimate. See Figure 1.3's notes for details on data, sample and income definitions.




Figure E.14. Department-Level Unemployment Rate in 1990



Figure E.15. Geographic Mobility by Parent Household Wage Rank

Notes: This figure presents the percentage of movers by parent income rank. Movers are defined as individuals whose adulthood department of residence is different from that of their childhood. See Figure 1.3 and 1.10's notes for details on data, sample and income definitions.




Figure E.16. Intergenerational Mobility and Geographic Mobility - Department Ranks

Notes: This figure represents the conditional expectation function of child household income rank with respect to parent household wage rank separately for individuals whose adulthood department of residence is different or not from their childhood department of residence. Percentile ranks are computed according to the local department income distribution. Parents are ranked within their department of residence in 1990 while children are ranked within their adulthood department. See Figures 1.3 and 1.10's notes for details on data, sample and income definitions.

F. Additional Tables

Birth Cohort	Born in Metropolitan France	+ Live with parents in 1990 census	+ At least one obs. in tax returns data (each inc. def.)	+ At least one obs. in tax returns data at 35-45	+ No parent in occupation 1 or 31
1972	9,083	7,946	7,515	7,515	7,015
1973	8,647	7,670	7,263	7,263	6,726
1974	8,704	7,713	7,294	7,294	6,758
1975	7,334	6,565	6,230	6,230	5,818
1976	7,762	6,963	6,567	6,547	6,100
1977	7,972	7,175	6,823	6,763	6,319
1978	7,755	7,000	6,691	6,585	6,136
1979	8,473	7,620	7,280	7,102	6,644
1980	8,822	7,965	7,559	7,239	6,774
1981	8,457	7,631	7,267	6,716	6,304
1972-1981	83,009	74,248	70,489	69,254	64,594

Notes: This table displays the number of observation for each child birth year cohort and the entire sample, at each sample restriction. Note that since parent income cannot be predicted for 23 children because one of their parents have an occupation not represented in the sample of synthetic parents of the corresponding gender, the actual sample size on which estimates are computed when using parent household wage as the parent income definition is 64,571 (i.e., 64,594 - 23).

Table F.1: Child Sample Construction

Characteristic	Synthetic Parents	Actual Parents
Females	53.42%	52.26%
Age in 1990	41.22%	40.74%
Born French	89.95%	88.36%
<i>1-digit occupation</i>		
1. Farmers	3.72%	3.47%
2. Craftsmen, salespeople, and heads of businesses	6.98%	6.77%
3. Managerial and professional occupations	9.76%	9.35%
4. Intermediate professions	15.48%	15.35%
5. Employees	20.76%	20.39%
6. Blue collar workers	23.19%	24.6%
7. Retirees	1.30%	1.32%
8. Other with no professional activity	18.81%	18.76%
<i>Education</i>		
No diploma	22.45%	23.8%
Primary education	19.38%	18.93%
BEPC	7.99%	8.18%
CAP	20.76%	19.91%
BEP	4.95%	5.00%
High school diploma	11.64%	11.47%
Bachelor or technical degree	6.08%	6.18%
Masters or PhD	6.75%	6.52%
<i>Country of birth</i>		
France	86.18%	84.81%
Maghreb	6.62%	8.03%
Other Africa	0.55%	0.73%
South Europe	3.32%	3.33%
Other Europe	2.33%	2.17%
Rest of the world	1.00%	0.94%
<i>Family structure</i>		
Single fathers	0.93%	0.72%
Single mothers	5.58%	5.25%
Both spouses active	58.73%	58.28%
Mother inactive	31.35%	32.32%
Father inactive	1.38%	1.38%
Both spouses inactive	2.03%	2.06%
<i>Municipality characteristics</i>		
Log population	7.83	7.85
Log density	0.46	0.49
% foreigners	2.31%	2.33%
Unemployment rate	6.22%	6.25%
% single mothers	6.3%	6.4%
N	134,572	140,136

Notes: See Section 3.2 for details on construction of samples. These statistics are computed before applying any income observation restrictions.

Table F.2: Average Characteristics of Actual and Synthetic Parents

2-digit occupation	Synthetic Parents	Actual Parents
Farmers with small farm	0.92%	0.84%
Farmers with medium-sized farm	1.22%	1.19%
Farmers with large farm	1.58%	1.44%
Craftsmen	3.62%	3.57%
Trade workers and related	2.62%	2.50%
Heads of company with ≥ 10 employees	0.73%	0.70%
Liberal profession	1.38%	1.32%
Public sector executives	1.07%	1.05%
Professors and scientific professions	2.12%	1.97%
Information, arts, and entertainment professions	0.32%	0.31%
Administrative executives and sales representatives	2.72%	2.66%
Engineers, technical executives	2.16%	2.05%
Teachers and related	2.64%	2.57%
Intermediate health and social work professions	2.48%	2.62%
Clerk, religious	0.01%	0.01%
Intermediate administrative professions of the public sector	1.54%	1.41%
Intermediate administrative professions and salesmen	4.06%	4.03%
Technicians	2.30%	2.29%
Foremen, supervisors	2.44%	2.42%
Civil servants	6.74%	6.69%
Police and military officers	1.27%	1.35%
Company administrative employees	6.92%	6.70%
Trade employees	2.24%	2.16%
Personal service workers	3.58%	3.49%
Skilled industrial workers	5.82%	6.14%
Skilled crafts workers	4.60%	4.83%
Drivers	2.19%	2.39%
Skilled handling, storing and transport workers	1.41%	1.47%
Unskilled industrial workers	6.19%	6.67%
Unskilled crafts workers	2.32%	2.42%
Agricultural workers	0.66%	0.69%
Former farmers	0.09%	0.07%
Former craftsmen, salespeople, and heads of businesses	0.10%	0.08%
Former managerial and professional occupation	0.09%	0.10%
Former intermediate professions	0.19%	0.17%
Former employees	0.33%	0.30%
Former blue collar workers	0.51%	0.60%
Unemployed who have never worked	0.36%	0.38%
Military contingent	0.00%	0.00%
Students ≥ 15 yrs old	0.10%	0.04%
Other inactive ≤ 60 yrs old	18.24%	18.20%
Other inactive ≥ 60 yrs old	0.10%	0.12%
N	134, 572	140, 136

Notes: See Table F.2's notes for sample construction.

Table F.3: Share of Actual and Synthetic Parents by 2-Digit Occupation

Gender	At least one child born in Metrop. France 1972-1981	+ Observed in 1990 Census	+ Born even year	+ At least two obs. at 35-45 in All Employee Panel	+ Not in occupation 1 or 31
Fathers	49,746	43,851	22,227	16,699	16,450
Mothers	52,904	48,097	24,297	15,104	14,973
All	102,650	91,948	46,524	31,803	31,423

Table F.4: Synthetic Parents Sample Construction

Child income definition	Parent income definition	Number of observations	0 child incomes (N.)	0 child incomes (%)	Negative child incomes (N.)	Negative child incomes (%)
Household income	Family income	64,571	0	0	41	0.06
Household income	Father income	57,902	0	0	35	0.06
Household wage	Family income	64,571	1976	3.06	0	0
Household wage	Father income	57,902	1690	2.92	0	0
Individual income	Family income	64,571	2479	3.84	68	0.11
Individual income	Father income	57,902	2162	3.73	60	0.1
Labor income	Family income	64,571	4990	7.73	0	0
Labor income	Father income	57,902	4376	7.56	0	0

Table F.5: Number of Observations by Child and Parent Income Definitions

	N	Missing (%)	Mean	Std. Dev.	25 th pctile	Median	75 th pctile
Synthetic Parents							
Synthetic father income (35-45 yrs old)	16,450	0	25,902	17,265	16,251	21,966	30,427
Number of syn. father income observations	16,450	0	7.66	2.42	6	8	9
Synthetic mother income (35-45 yrs old)	14,973	0	15,167	10,143	7,496	14,140	21,027
Number of syn. mother income observations	14,973	0	6.95	2.84	5	7	9
Parents							
Fraction single parents in 1990	11.72%						
Fraction female among single parents	88.3%						
Father age at child's birth	64,594	10.35	28.48	6.08	24	28	31
Mother age at child's birth	64,594	1.37	25.89	5.15	22	25	29
Father age in 1990	64,594	10.35	41.98	6.61	38	41	45
Mother age in 1990	64,594	1.37	39.42	5.81	35	39	43
Children							
Household income (average 2010-16)	64,594	0	46,599	38,371	27,696	41,417	56,481
Household wage (average 2010-16)	64,594	0	38,460	30,184	20,812	35,205	50,096
Individual income (average 2010-16)	64,594	0	23,512	20,471	14,375	21,159	28,737
Labor income (average 2010-16)	64,594	0	21,092	19,120	10,067	19,877	27,487
Fraction female	49.77%						

Notes: See Sections 3.2 and 3.3 for details on sample construction and income definitions.

Table F.6: Descriptive Statistics

	Intergenerational Elasticity	First-Stage Instruments	Data	Income Definitions	Child Age
Lefranc and Trannoy (2005)	0.4-0.438 ¹	Education (8 cat.) + occupation (7 cat.)	FQP	labor earnings (excl. UI) ²	30-40
Lefranc (2018)	0.577 ³	Education (6 cat.)	FQP	labor earnings (excl. UI) ²	28-32
EqualChances.org	0.357	Education (3 cat.) + occupation (9 cat.)	Synthetic fathers: ECHP Sons: EU-SILC	net personal employee income	-

Our estimate	0.443				

Notes: FQP = Formation-Qualification-Profession; ECHP = European Community Household Panel; EU-SILC = European Union Statistics on Income and Living Conditions

¹ Estimates taken from Table I, Panel A, cols. (1)-(4), p.65.

² Only salaried workers.

³ Estimates taken from Table 2, 1971-75, col. (2), p.823.

Table F.7: Comparison with Existing Father-Son IGE Estimates for France

	IGE	RRC	AUM
First-stage MSE	-0.160 (0.127)	-0.088 (0.056)	1.400 (3.487)
Constant	0.565*** (0.053)	0.318*** (0.024)	42.370*** (1.465)
Observations	85	85	85
R ²	0.019	0.029	0.002

Notes: Standard errors in parentheses. *p<0.1; **p<0.05; ***p<0.01.

Table F.8: Department-Level MSEs and Measures of Intergenerational Income Mobility

	Department	Observations	IGE	RRC	AUM
01	Ain	535	0.4	0.27	47.5
02	Aisne	735	0.63	0.4	39
03	Allier	365	0.45	0.23	42.9
04	Alpes-de-Haute-Provence	141	*	*	*
05	Hautes-Alpes	112	*	*	*
06	Alpes-Maritimes	773	0.37	0.24	43.6
07	Ardèche	313	0.4	0.19	42.4
08	Ardennes	376	0.56	0.31	38.8
09	Ariège	121	*	*	*
10	Aube	361	0.27	0.2	43.2
11	Aude	274	0.77	0.38	36.9
12	Aveyron	243	0.37	0.25	44.5
13	Bouches-du-Rhône	1,795	0.62	0.33	41.9
14	Calvados	781	0.47	0.31	42.7
15	Cantal	164	*	*	*
16	Charente	374	0.52	0.24	40.6
17	Charente-Maritime	559	0.45	0.3	40.9
18	Cher	370	0.48	0.26	42.9
19	Corrèze	219	0.56	0.38	40.9
20	Corse	236	0.48	0.17	45.6
21	Côte-d'Or	549	0.44	0.29	44.1
22	Côtes-d'Armor	590	0.3	0.25	44.6
23	Creuse	102	*	*	*
24	Dordogne	337	0.31	0.25	39.9
25	Doubs	635	0.33	0.25	47.6
26	Drôme	435	0.52	0.33	39
27	Eure	738	0.42	0.25	43.9
28	Eure-et-Loire	505	0.53	0.31	43.1
29	Finistère	979	0.44	0.22	44.7
30	Gard	577	0.63	0.28	39.3
31	Haute-Garonne	949	0.51	0.29	43.4
32	Gers	136	*	*	*
33	Gironde	1,304	0.43	0.25	41.8
34	Hérault	788	0.65	0.32	38.3
35	Ille-et-Vilaine	1,036	0.36	0.26	43.6
36	Indre	235	0.74	0.39	37
37	Indre-et-Loire	597	0.63	0.31	41.8
38	Isère	1,217	0.43	0.25	43.5
39	Jura	269	0.37	0.27	48.7
40	Landes	326	0.39	0.22	43
41	Loir-et-Cher	357	0.57	0.25	43.2
42	Loire	901	0.39	0.27	44.5
43	Haute-Loire	194	*	*	*
44	Loire-Atlantique	1,467	0.48	0.26	42.8

Notes: * Insufficient number of observations (< 200).

Table F.9: Department-Level Intergenerational Mobility Estimates

	Department	Observations	IGE	RRC	AUM
45	Loiret	706	0.63	0.36	41.2
46	Lot	137	*	*	*
47	Lot-et-Garonne	319	0.74	0.27	40.4
48	Lozère	63	*	*	*
49	Maine-et-Loire	931	0.49	0.31	40.8
50	Manche	566	0.56	0.29	43.3
51	Marne	676	0.37	0.33	43.1
52	Haute-Marne	263	0.62	0.35	41
53	Mayenne	329	0.67	0.34	40.6
54	Meurthe-et-Moselle	862	0.59	0.32	42.6
55	Meuse	238	0.32	0.24	44.9
56	Morbihan	778	0.45	0.27	43.3
57	Moselle	1,274	0.52	0.31	44.7
58	Nièvre	251	0.31	0.18	43
59	Nord	3,668	0.63	0.36	38.7
60	Oise	1,008	0.42	0.24	44.1
61	Orne	357	0.65	0.33	38.9
62	Pas-de-Calais	2,145	0.7	0.39	36.8
63	Puy-de-Dôme	664	0.41	0.26	43.6
64	Pyrénées-Atlantiques	571	0.49	0.28	43.5
65	Hautes-Pyrénées	209	0.48	0.21	41.4
66	Pyrénées-Orientales	356	0.75	0.29	38.1
67	Bas-Rhin	1,033	0.54	0.31	45.2
68	Haut-Rhin	792	0.53	0.29	46.7
69	Rhône	1,583	0.46	0.27	45.2
70	Haute-Saône	273	0.88	0.28	40.7
71	Saône-et-Loire	661	0.69	0.39	42.4
72	Sarthe	635	0.57	0.36	40
73	Savoie	430	0.45	0.26	45.7
74	Haute-Savoie	629	0.4	0.19	54.4
75	Paris	1,352	0.51	0.28	49.8
76	Seine-Maritime	1,547	0.54	0.34	41.9
77	Seine-et-Marne	1,529	0.4	0.24	46.2
78	Yvelines	1,645	0.47	0.25	49.2
79	Deux-Sèvres	376	0.35	0.27	43.2
80	Somme	737	0.45	0.3	40.4
81	Tarn	354	0.54	0.36	38.2
82	Tarn-et-Garonne	202	0.59	0.23	40.8
83	Var	773	0.59	0.31	39
84	Vaucluse	468	0.5	0.27	42.6
85	Vendée	627	0.37	0.19	44.6
86	Vienne	464	0.51	0.33	41.7
87	Haute-Vienne	357	0.56	0.36	40.5
88	Vosges	504	0.46	0.26	40.6

Notes: * Insufficient number of observations (< 200).

Table F.9: Department-Level Intergenerational Mobility Estimates (*continued*)

	Department	Observations	IGE	RRC	AUM
89	Yonne	388	0.32	0.23	43
90	Territoire de Belfort	172	*	*	*
91	Essonne	1,302	0.47	0.28	48
92	Hauts-de-Seine	1,248	0.5	0.28	48.9
93	Seine-Saint-Denis	1,495	0.45	0.2	47.3
94	Val-de-Marne	1,188	0.41	0.19	51
95	Val-d'Oise	1,366	0.49	0.26	46.7

Notes: * Insufficient number of observations (< 200).

Table F.9: Department-Level Intergenerational Mobility Estimates (*continued*)

Child income definition	IGE-RRC	RRC-AUM	IGE-AUM
Household income	0.65	-0.57	-0.55
Individual income	0.72	-0.55	-0.45
Individual wage	0.70	-0.41	-0.26

Notes: See Figure 1.8 for corresponding maps.

Table F.10: Correlation Between Department-Level Intergenerational Mobility Measures

	<i>Dependent variable: Child household income rank</i>				
	(1)	(2)	(3)	(4)	(5)
Parents income rank	0.259*** (0.005)	0.259*** (0.005)	0.258*** (0.005)	0.162*** (0.007)	0.135*** (0.012)
Mover ($\hat{\gamma}$)	4.572*** (0.472)	4.591*** (0.471)	4.897*** (0.478)	4.926*** (0.477)	4.883*** (0.478)
Parents income rank \times Mover ($\hat{\delta}$)	-0.014* (0.008)	-0.014* (0.008)	-0.016** (0.008)	-0.026*** (0.008)	-0.027*** (0.008)
Constant	36.401*** (0.265)	36.137*** (0.279)	35.574*** (1.125)	26.815*** (1.570)	28.162*** (1.620)
Birth cohort	✓	✓	✓	✓	✓
Gender		✓	✓	✓	✓
Department FE			✓	✓	✓
Parents' education				✓	✓
Parents' 2-digit occupation					✓
$\mathbb{E}[\hat{\gamma} + \hat{\delta}p_{p,i}] = \hat{\gamma} + \hat{\delta} \times 50.5$	3.87	3.88	4.09	3.61	3.52
$\mathbb{E}[\hat{\gamma} + \hat{\delta}p_p p_p = 25]$	4.22	4.24	4.50	4.28	4.21
$\mathbb{E}[\hat{\gamma} + \hat{\delta}p_p p_p = 75]$	3.52	3.54	3.70	2.98	2.86
Observations	64,571	64,571	64,571	64,571	64,571
Adjusted R ²	0.074	0.074	0.077	0.089	0.095

Notes: Bootstrapped standard errors in parentheses. *p<0.1; **p<0.05; ***p<0.01.

Table F.11: Intergenerational & Geographic Mobility - Department Ranks

	<i>Dependent variable: Child household income rank</i>				
	(1)	(2)	(3)	(4)	(5)
Parents income rank	0.278*** (0.005)	0.278*** (0.005)	0.271*** (0.005)	0.171*** (0.008)	0.153*** (0.017)
Destination department (ref.: stayers)					
Low-income	0.902 (0.618)	0.923 (0.618)	1.046* (0.616)	0.852 (0.616)	0.685 (0.616)
Medium-income	11.355*** (0.951)	11.373*** (0.952)	10.846*** (0.945)	11.045*** (0.948)	11.027*** (0.950)
High-income	18.819*** (1.224)	18.839*** (1.224)	18.265*** (1.247)	18.465*** (1.260)	18.567*** (1.258)
Parents income rank × Low-income	-0.019* (0.011)	-0.019* (0.011)	-0.017 (0.011)	-0.020* (0.011)	-0.018 (0.011)
Parents income rank × Medium-income inc	-0.042*** (0.013)	-0.042*** (0.013)	-0.038*** (0.013)	-0.051*** (0.013)	-0.052*** (0.013)
Parents income rank × High-income	-0.035** (0.016)	-0.035** (0.016)	-0.035** (0.016)	-0.054*** (0.016)	-0.058*** (0.016)
Constant	34.143*** (0.261)	33.860*** (0.277)	37.460*** (1.213)	28.392*** (1.655)	29.369*** (1.779)
Birth cohort	✓	✓	✓	✓	✓
Gender		✓	✓	✓	✓
Department FE			✓	✓	✓
Parents' education				✓	✓
Parents' 2-digit occupation					✓
Observations	64,571	64,571	64,571	64,571	64,571
Adjusted R ²	0.118	0.118	0.124	0.135	0.142

Notes: Bootstrapped standard errors in parentheses. *p<0.1; **p<0.05; ***p<0.01.

Table F.12: Intergenerational Mobility and Income Level in the Destination Department

2

Older Schoolmate Spillovers on Higher Education Choices

This chapter is based on a paper co-authored with Nagui Bechichi (PSE).

Abstract

This paper studies within-high school spillovers in college and major choices. Specifically, we examine how students' higher education choices are influenced by the higher education trajectory of older schoolmates. We exploit admission cutoffs generated by the French centralised admission system and compare high schools with a marginally admitted student to a college-major to high schools with a marginally rejected student. Using exhaustive administrative application data, we find that students are about 7 percentage points (+25%) more likely to apply to the same college-major as a marginally enrolled older schoolmate, and 3 percentage points (+67%) more likely to enrol. These 2SLS estimates correspond roughly to 60% of the magnitude of sibling spillovers estimated by [Altmejd et al. \(2021\)](#). We explore two potential mediating factors for these within-high school spillovers: (i) teachers, and (ii) student role models. Our early evidence on these mechanisms suggests role model effects play the dominant role. Students with the same "principal" teacher as the previous cohort's marginally admitted student are not more likely to apply to or enrol in the same college-major compared to other students. In contrast, students with the same gender or same socio-economic status as the older schoolmate are significantly more likely to apply to and enrol in the same college-major. These results highlight the important role played by students' high school environment in shaping their higher education choices.

1. Introduction

HOW do students choose whether and where to apply to university? Considering the large returns to higher education, and the large differences across majors and institutions, this question has received tremendous attention. Recent work has highlighted the important roles played by informational deficits ([Hoxby and Turner,](#)

2013a; Carrell and Sacerdote, 2017), and by students' social networks, such as their parents (Altmejd, 2023), their siblings (Aguirre and Matta, 2021; Altmejd et al., 2021) and even their neighbors (Barrios-Fernández, 2022). This suggests exposure to peers' higher education choices might be an important source of information for students' decisions. However, we know very little about how high school peers shape this decision. In particular, within the same high school, how are students' applications and enrollment choices influenced by older schoolmates' higher education trajectories? Causally identifying such effects is challenging due to the important sorting of students across high schools and the endogeneity of students' higher education choices to their high schools.

This paper provides the first causal evidence on within-high school¹ spillovers on higher education choices. Using administrative application data from France covering close to 90% of higher education programs between 2013 and 2017 (Bechichi et al., 2021), we show that students are more likely to apply to and enrol in a college-major² if a student from the same high school enrolled in this exact same college-major the previous year. We also find important spillovers on the choice of college more broadly, but no effects on the chosen major.

We identify within-high school spillovers by exploiting admission cutoffs generated by France's centralised admission procedure. This allocation ensures programs cannot anticipate ex-ante the high school of the last admitted student. As such high schools around a college-major's admission threshold are virtually identical other than for having a student ranked just above or just below the rank of the last admitted student to this college-major. This generates quasi-random variations in the college-majors to which a high school's students are admitted to and enrol in, which in turn also generates quasi-random variations in the college-majors to which the following cohort of students in the same high school is exposed to. This enables us to implement a regression discontinuity design to estimate within-high school spillovers on applications and enrollment. While existing work has exploited comparable cutoffs generated by *academic* thresholds in admission policies (e.g., Altmejd et al. (2021); Estrada et al. (2022)), our design is very similar in spirit except we only observe the relative *ranking* of students by the college-majors to which they have applied. Since several students from the same high school may apply to the same college-major, we keep only the high school's best ranked applicant by the college-major, as in Estrada et al. (2022).

¹Technically, our analysis is undertaken at the high school x track level because, in France, high school tracks are very segregated within high schools and higher education programs are often largely track-specific. To ease legibility, we use "high school" to refer to "high school x track".

²We will use interchangeably *college-major*, *college program* and *program*.

We find that students follow the higher education choices of their high school's previous graduating cohort. They are 7 percentage points (+25% relative to the counterfactual mean) more likely to apply to, and 3 percentage points (+67%) more likely to enroll in, a college-major in which a student from their high school's previous cohort was marginally admitted to and enrolled in, relative to students in high schools with a marginally rejected older schoolmate. We also uncover large impacts on the intensive margin, i.e., the number, of applications and enrolled students: 0.24 (+30%) and 0.05 (+72%) percentage points increases respectively.

The magnitude of these effects is large. Compared to sibling spillovers in college-major estimated by [Altmejd et al. \(2021\)](#) for Chile, Croatia and Sweden, our within-high school spillovers are between 43% and 78% as large as the impact on applications they find, and between 40% and 88% of their enrollment effects. Moreover, we also show that students are more likely to apply to and enrol in the same college as their older high school peers but there are no spillovers on major choice. The (relative) magnitude of the college spillovers are roughly similar to the college-major spillovers, 9.6 percentage points (+17%) for applications and 10.9 percentage points (+52%) for enrollments. The lack of spillovers on majors could potentially be explained by the fact that students have stronger preferences over what they want to study than over where they want to study or because they are more aware of existing majors. Therefore the college component of a previous peer's enrollment is more salient to them. This result is in line with [Altmejd et al. \(2021\)](#) and [Aguirre and Matta \(2021\)](#) who also find no sibling spillovers on majors.

We uncover several insightful heterogeneities with respect to college-major spillovers. First, we find that the magnitude of the spillovers are broadly constant over the four outcome years. This suggests they are not the result of a given year's idiosyncrasies, but rather a structural determinant of students' higher education choices. Second, with regards to student characteristics, we find that within-high school spillovers on applications are of similar magnitude for both genders, though the effects on enrollment are significantly larger for boys. This could be driven by differences in the types of degrees applied to. Moreover, and quite surprisingly, we find that low socioeconomic status (based on legal guardian's occupation) students are only slightly more responsive than their very high SES peers. This is somewhat unexpected since, a priori, one might suppose very high SES students to be better informed about higher education and thus not be much influenced by older schoolmates' higher education trajectory. Conversely, low SES students tend to be less aware of the higher education landscape ([Hoxby and Turner, 2013b](#)) and thus one could expect they would be more influenced by peers.

Third, the magnitude of spillovers vary across some characteristics of high schools. All high school tracks display spillovers of roughly the same magnitude, with slightly larger ones (in percentage terms) for the literature track. This result is quite noteworthy because it suggests that the acquisition of information about higher education choices is relevant in very different contexts. That being said, spillovers are largest in small high schools (less than 30 students), and are decreasing in high school size. This could be the result of several explanations. Small high schools may exhibit closer relationships between students and their teachers, and thus teachers might be better aware of their students' higher education choices. As such they may encourage their subsequent classes to apply to the same college-majors as their past students. Another explanation could be that smaller high schools maintain better links with their alumni through, for example, annual alumni gatherings. A last explanation could simply be that smaller high schools are located in more rural areas where information about college-majors may be more scarce and therefore informational shocks are amplified to a much larger extent than in information-rich high schools located in large cities. Moreover, we find that high schools in the second and top quintile in the high school academic level distribution (measured as the median of its students' end of high school exam grades) exhibit the largest spillovers on applications. It is not clear exactly what could explain this result.

Fourth, we explore how spillovers differ across college-major characteristics. We find that spillovers are largest for public university, technical and vocational programs, but not for preparatory classes, which tend to be quite prestigious, nor for other types of college-majors. In line with these results, college-majors in the bottom 10% of selectivity (proxied by the median end-of-high school exam grade of enrolled students) exhibit the largest spillovers, and they are decreasing in selectivity. There are no spillovers for college-majors in the top 10%. This is somewhat surprising, as one may expect very selective college-majors to be those where some students may not dare to apply. This leads us to infer that students are learning about college-majors that they had been unaware of before rather than increasing their confidence to apply to prestigious college-majors.

Lastly, we assess how the interaction between high schools' and college-majors' characteristics may shape within-high school spillovers. The results suggest geographically close (less than 25 km) and moderately far (between 50 and 100 km) programs induce the largest spillovers. In terms of the intensive margin of applications and enrollment these moderately far programs display significant cross-cohort spillovers. Additionally, we find that students in low-achieving (bottom 25%) high schools are significantly more likely to follow an older schoolmate marginally admitted to a college-major in the top

25% of selectivity. This effect could be interpreted as raising the aspirations or awareness of these high schools' top performing students. Conversely, we find intriguingly that spillovers are quite large for top performing (top 25%) high schools for whom the marginally admitted student went to a college-major in the bottom quartile of selectivity.

In the last section of the paper we explore two mechanisms that may underpin our within-high school spillovers: (i) the role of teachers, and (ii) student role model effects. These mechanisms are not mutually exclusive but they lead to drastically different policy recommendations. We estimate the extent to which cross-cohort spillovers may be driven by teachers, for example, by recommending their past students' college-majors. Since we do not directly observe all of students' teachers, we test this mechanism in two complementary ways. First, we examine whether students are more likely to follow an older schoolmate if they share the same "*principal*" teacher. In France, each class is assigned a principal teacher who is in charge of the class' administrative duties over the course of the academic year, and in particular, helps and supervises students' higher education applications. Second, we assess whether students are more likely to follow an older schoolmate if they are in the same class identifier (e.g., senior class A, senior class B), an imperfect proxy for sharing the same set of teachers as that older schoolmate. We find that students sharing the same principal teacher or the same class identifier are equally likely to follow the marginally admitted older schoolmate's higher education choices as students with different teachers or principal teacher. This appears to suggest a rather limited direct role for teachers, at least in explaining the within-high school spillovers we document. This could be because teachers help their students by recommending a wide range of college-majors rather than only those of their past students.

Second, we attempt to disentangle whether our spillovers are more likely due to informational shocks or to role models effects. To test this, we assess whether the effects are larger for students sharing the same gender or socio-economic status as the marginally admitted older schoolmate. We interpret this test as capturing a role model effect rather than an information effect since, a priori, the marginally admitted student's gender or SES does not affect the informational content of his or her higher education trajectory but affects the way this information is perceived. We find strong evidence in favor of role model effects. Girls are significantly more likely to apply to college programs when the marginally admitted older schoolmate was a girl (+9%) but not when it was a boy (+3%, insignificant), while boys are more likely to follow a boy (+8%) but not a girl (+2%, insignificant). Similarly, low SES students are significantly more likely to apply to a degree when the marginally admitted older schoolmate was also of low SES background (+13%),

but not when the latter is from a very high SES background (+1%, insignificant). However, very high SES students are largely unresponsive regardless of the SES of the treated older schoolmate. This is consistent with them being more knowledgeable about or having stronger preferences for college-majors.

Our paper contributes to several strands of literatures. First, and most directly, it contributes to a very small literature on spillovers across school cohorts. To the best of our knowledge, the only other paper focusing on such spillover effects is [Estrada et al. \(2022\)](#). They find that the admission of an older schoolmate into an elite high school in Chile significantly increases the number of applications and admissions to the same elite high school the following year. We expand on this work by providing evidence of similar within-school spillovers for (close to) the universe of college-majors in France and by exploring potential underlying mechanisms. We believe such spillovers could be observed in other educational settings such as the choice of graduate school or in the labor market such as the choice of first employer upon graduating (e.g., [Kramarz and Skans \(2014\)](#)).

Second, our paper contributes to the literature on peer effects in education, and higher education specifically (see [Barrios-Fernandez \(2023\)](#) for a nice overview of these literatures). The closest papers relate to the role played by students' family members (parents and siblings) and close peers (e.g., neighbors, teachers, or classmates) in shaping their higher education application behavior. Several papers have highlighted the large causal impact of such relations. Using Swedish application data, [Altmejd \(2023\)](#) finds that children are 50% more likely to graduate from the same major as their parents. Additionally, [Aguirre and Matta \(2021\)](#) and [Altmejd et al. \(2021\)](#) document the very large role played by older siblings' college-major enrollment on their younger siblings' higher education choices in Chile, Croatia, Sweden and the United States. They find no effect of siblings on major choice, contrary to [Avdeev et al. \(2023\)](#) who find spillovers on major choice in the Netherlands. Moreover, [Barrios-Fernández \(2022\)](#) reveals that having a close neighbor at the margin of enrolling in higher education increases one's likelihood to enrol in higher education. Collectively, these studies highlight the predominant role of students' social ties in influencing their higher education choices. Our paper adds to this literature by analysing a previously overlooked dimension of students' social network, their older schoolmates.

Third, our paper contributes to the literature on role models in educational settings. This literature emphasises the important influence same-gender ([Carrell et al., 2010](#); [Canaan and Mouganie, 2021](#)) and same-race teachers ([Gershenson et al., 2022](#)) have on their stu-

dents' higher education outcomes and major choice. Other studies also find evidence that external role models, such as successful female scientists, can steer young girls towards scientific higher education tracks (Breda et al., 2023). We add to this literature by highlighting a new setting in which such effects operate, older high school peers, as well as by documenting that another characteristic of students, their socioeconomic background, can be leveraged to induce role model effects.

Lastly, this paper contributes to the broad literature on the determinants of college and major choice. Recent work has emphasized vast differences in returns across majors and to a more moderate extent across higher education institutions, underscoring how important both choices are to individuals' future outcomes (Altonji et al., 2012; Hastings et al., 2013; Kirkeboen et al., 2016; Aucejo et al., 2022). Our paper sheds light on determinants that occur in students' high schools. Other papers have highlighted the long-term consequences of teachers (Chetty et al., 2011, 2014; Jackson, 2018), and in particular on higher education attendance. Moreover, recent work by Mulhern (2023) points out the role played by high school guidance counselors. We stress that high schools not only influence college-going behavior but also play a role in shaping students' specific college-major choices.

The rest of the article is organized as follows. Section 2 describes the institutional background and data. The empirical strategy is presented in Section 3, while Section 4 discusses the main results, their robustness and some insightful heterogeneity. In Section 5, we investigate the mechanisms that might explain our within-high school cohort spillovers. Section 6 concludes.

2. Institutional Background and Data

This section describes the most important features of the French higher education system and the data we use. We provide additional background details in Appendix A. Central to our analysis, students in France apply to a college-major combination through a centralised application platform. There are no nationally set criteria for admission, each college-major decides on its own (undisclosed) criteria based on all the available information (academic and non-academic) in students' application package. We also discuss how the allocation algorithm generates the discontinuities we exploit to identify within-

high school spillovers. Note that we undertake our analysis at the *high school x track*³ level rather than at the *high school* level. This is because (i) in France, high school tracks are very segregated within high schools, and (ii) many college-majors are largely restricted to specific high school tracks. To ease legibility, we continue referring to high schools even though technically we mean high school x track. Thus all statistics related to high schools are computed at the high school x track level.

2.1. Institutional Background

From 2009 to 2017, senior high school students in France applied to college-majors through a centralised online application and admission platform called *Admission Post-Bac*.⁴ Over the period, this platform progressively covered up to 90% of higher education programs (Bechichi et al., 2021), with over 10,000 programs and 800,000 applicants in 2017.⁵ The allocation of students to college-majors was spread over three stages. First, students would submit a rank-ordered list of programs with up to 36 choices. Second, programs would rank their applicants and submit their capacity constraints (i.e., how many available seats they have) to the platform. Lastly, offers were sent out following a three-round college-proposing deferred acceptance algorithm.⁶ The three rounds were to account for seats freeing up due to (i) students choosing programs outside the platform or directly entering the labor market, or (ii) students failing the end of high school exam, the *Baccalauréat*, which is necessary to enrol in higher education (see Appendix Figure A.2 for a timeline of the entire procedure).

³There are 13 high school tracks, grouped into three aggregate tracks (general, technological, and professional). The three general tracks are Sciences (S), Social Sciences (ES), and Literature (L). The 8 technological tracks are Management Sciences and Technologies (STMG), Sustainable Development Sciences and Technologies (STI2D), Health Sciences and Technologies (ST2S), Laboratory Sciences and Technologies (STL), Design and Applied Arts Sciences and Technologies (STD2A), Agronomy and Life Sciences and Technologies (STAV), Hospitality (H), and Music and Dance Techniques (TMD). The two broad professional tracks are Professional (P) and Agricultural Professional (PA).

⁴This platform has since been replaced by *Parcoursup*, which changed a number of the parameters of the application and admission system.

⁵Various schools such as Paris Dauphine, the Institutes of Political Studies (IEP), some private programs, and nursing schools were not on the platform.

⁶The college-proposing deferred acceptance algorithm implemented is actually a slight variant from the Gale Shapley algorithm to account for joint applications to a college-major and the college-major's student accommodation. This deviation is irrelevant for our study. See Charousset et al. (2021) Appendix 2.A.2. for additional details.

2.2. Data

We use application-level administrative data from the *Admission Post Bac* application platform for the years 2013-2017. They contain comprehensive information on students' rank-ordered list of applications to college-majors, college-majors' rankings of applicants, and the matching outcome. We do not observe enrollment per se, we only observe whether students have *accepted* an offer from a college-major.⁷ Some students may accept an offer but end up not enrolling. Importantly, we also have detailed information on students' background characteristics such as their high school, their high school track, their end of high school exam grades, and their socioeconomic status based on their legal guardian's occupation. High schools and programs can be tracked over time through a unique identifier.

3. Empirical Strategy

Our aim is to identify the extent to which students apply to and enrol in the same college-major in which a student from their high school's previous graduating cohort enrolled. The idea is that this student's experience at the college-major could spillover to younger cohorts of students, through some form of increased awareness of or information about this college-major. Identifying such within-high school spillovers is challenging because students across cohorts, by definition, share the same high school characteristics such as teachers, geographic location, and general resources, which all have independent effects on higher education choices regardless of older schoolmates' higher education trajectories.

For a given college-major, comparing all high schools with an admitted student to all high schools with a rejected student would likely overstate cross-cohort spillovers because high schools with admitted students are very likely different from those with rejected students. Indeed they may have more applicants to this college-major, have more admitted students, or have different academic achievement levels, which could correlate with the likelihood of applying. College-majors could also favor high schools from which they are used to admitting students.

We use admission rank cutoffs to overcome these challenges and identify within-high

⁷We are currently working on matching the enrollment data with the application data. The technical difficulty hinges on the fact that though there are common student identifiers in both datasets, there are no common college-major identifiers. This implies generating a one-to-one mapping between the program identifiers in both datasets.

school spillovers in higher education choices in a (fuzzy) regression discontinuity design. The fuzzy nature of the design stems from the fact that we define (conceptually) high schools as being treated if an older schoolmate *enrolled* in the college-major, not only if he or she was *admitted*. Unlike other countries where programs select applicants following clearly laid out academic criteria, in France, programs' rankings were based on undisclosed and unknown criteria using all the available information present in students' application package, such as their high school grades, their teachers' comments, and any other salient information such as their high school or geographic location (see [Charouset et al. \(2021\)](#) for an attempt to reverse engineer the potential criteria used). We can nonetheless exploit the cutoff rank of the last admitted student to a college-major generated by the admissions process. There is one cutoff for each college-major in each year.

These cutoffs induce a discontinuity in the likelihood that a high school has an older student who is admitted and enrolls in the college-major. The intuition behind our identification strategy is to compare high schools that are otherwise identical except for having an older student just above or just below the admissions cutoff for a given college-major. As such, the only difference between these high schools is that an older student has a greater probability of receiving an offer from the program (depending on the relative ranking of the program in their rank-ordered list) and eventually enrolling in the college-major and therefore influencing the high schools' younger cohorts' higher education choices. Students ranked below the last admitted, by definition, could not receive an offer.

These cutoff ranks cannot be manipulated by applicants or college-majors for two reasons. First, cutoffs depend on applicants' rank-ordered list, which are unknown to programs when reporting capacities and rankings. Second, the evolution from the first to the third admission cutoff is determined by applicants' decision to accept or decline offers and thus cannot be anticipated by neither students nor programs during the application phase. Since high schools on either side of the cutoff are essentially identical, this rules out the possibility that the results are driven by differences in high schools' student body composition, size or location. Moreover, we can rule out the reflection problem because the older cohort's enrollment decisions cannot be impacted by younger schoolmates' higher education choices ([Manski, 1993](#)).

3.1. *Running Variable*

Since several students from the same high school may apply to the same college-major, we follow [Estrada et al. \(2022\)](#), and only keep each high school's best ranked applicant by each college-major. As such for each high school, each college-major, in each year, we

only have one observation, i.e., that of the best ranked applicant from that high school to that college-major in that year. By taking the rank of a high school’s best ranked applicant by a given college-major, we ensure that we do not misclassify the high school’s treatment status with respect to the college-major, for example in cases where the high school has one rejected and one admitted student. We then center these high schools’ best ranks around the college-major’s rank of the last admitted student. Specifically, for each high school s with student(s) $i(s)$ applying to college-major j in year t , we define the running variable as:

$$\text{distance to last admitted student}_{s jt} = \text{rank of last admitted student}_{jt} - \max_{s jt} \{ \text{rank}_{i(s) jt} \}$$

If *distance to last admitted student*_{s jt} is greater than or equal to 0, then high school s is (potentially) treated by college-major j in year t in the sense that the best ranked applicant might have received an offer and enrolled in college-major j , otherwise the high school is not. Appendix B provides additional details on the running variable as well as an illustration of how the running variable is computed for one college-major in a given year. Note that because it is very common for college-majors to have several applicants from the same high school, by construction our running variable for a given college-major does not necessarily have actual observations at each rank.

3.2. Estimation Sample

We restrict our sample to (i) college-majors with at least one applicant ranked after the last admitted student (86% of all college-majors - year), (ii) high schools with at least one applicant in two consecutive years⁸ (97% of all high schools - year), and (iii) college-majors reporting the same capacity constraint over the three admission rounds (85% of all college-majors - year). Restriction (i) ensures that, for each college-major, there is (potentially) at least one high school with a rejected student, restriction (ii) ensures we can observe high schools’ outcomes in the treatment year, i.e., the year with a marginally admitted older schoolmate, and in the following year, and restriction (iii) ensures college-majors cannot attempt to manipulate who the last admitted student will be.

Moreover, we impose additional conditions to our sample to ascertain that high schools and college-majors are as similar as possible around the cutoff. We thus add the following restrictions: (i) we “symmetrize” the running variable such that, for example, if a college-

⁸Missing values correspond to cases where there are no applicants on the platform from the high school in the outcome year.)

major has n applicants ranked after the last admitted student, we keep n applicants to the right of the last admitted student, (ii) we drop the last admitted student and (iii) we keep only college-majors with at least two observations on both sides of the cutoff within the chosen regression discontinuity bandwidth (after symmetrization and dropping the last admitted student). Our final sample, within the bandwidth, contains 18,543 college-majors - year and 46,755 high schools - year, for a total number of observations of 354,168.

Table 2.1 presents some descriptive on high schools and college-majors in the raw sample⁹ (*Full sample*) and in the sample used for the analysis (*RD sample*). The full sample and analysis samples are actually very similar, on average, in terms of high school characteristics, be it size, student body composition and academic achievement. Our analysis sample slightly over-represents technological high schools tracks and slightly under-represents professional tracks. With respect to college-major characteristics, the analysis sample contains college-majors with more applicants, though they have very similar average academic records, and with more high schools within the application pool. Perhaps most importantly, our analysis contains more programs in public universities¹⁰ (31.5% in the analysis sample vs 24.2% in the full sample), and a bit less of the other types of programs (in particular preparatory classes, 5.6% vs 8.1%).

⁹Our raw sample is restricted to applications from high school students with a non-missing high school identifier.

¹⁰One reason why public universities degrees may be over-represented in our analysis sample is that these degrees rank *all* applicants while other degrees rank as many applicants as they wish.

Table 2.1: Descriptive Statistics

	Full sample <i>All high school applicants</i> 2013-2016 (1)	RD sample <i>Distance to cutoff rank</i> $\in [-21, 21]$ (2)
<i>Panel A. High school characteristics</i>		
Number of high schools	53,421	46,755
Mean number of students	41.1	45.8
Female (%)	53.7	54.2
Mean Bac grade	12.2	12.2
Highest honors at Bac (%)	6.2	6.4
Very high SES (%)	28.3	28.1
High SES (%)	13.3	13.7
Middle SES (%)	30.4	30.4
Low SES (%)	24.5	24.8
Missing SES (%)	3.5	3.0
Scientific high-school track (%)	19.3	20.5
Social science high-school track (%)	17.4	18.1
Literature high-school track (%)	15.2	15.1
Technological high-school track (%)	23.5	25.0
Professional high-school track (%)	24.5	21.4
<i>Panel B. College-major characteristics</i>		
Number of college-majors	40,844	18,543
Number of applicants	311.8	359.3
Mean Bac grade of applicants	12.2	12.2
Highest honors at Bac of applicants (%)	6.4	6.3
Number of admitted students	37.6	38.5
Mean Bac grade of admitted students	12.3	12.4
Highest honors at Bac of admitted students (%)	6.5	6.8
Number of high schools within application pool	123.0	143.7
Public university (%)	24.2	31.5
Vocational degree (STS) (%)	49.8	47.6
Technical degree (IUT) (%)	8.5	7.6
Academic preparatory classes (CPGE) (%)	8.1	5.6
Other institutions (%)	9.3	7.6
Number of observations	5,889,043	354,168

Notes: This table shows descriptive statistics for two samples: (1) the full sample, i.e., all high school applicants between 2013-16, and (2) our regression discontinuity (RD) sample. *High school* refers to high school x tracks - year, while *college-majors* refer to college-major - year. The Bac is the French end of high school exam.

3.3. Empirical Specification

First-stage. Before discussing our precise estimation, we illustrate the fuzzy nature of our regressions discontinuity design. Specifically, we use all college-majors in our sample and stack them such that they are all centered around the rank of the last admitted student. Figure 2.1 shows the likelihood that a high school has at least one enrolled student (panel A) and the number of enrolled students (panel B) as a function of the high school's distance to the last admitted student. The probability of treatment increases from essentially 0%¹¹ to the left of the cutoff to 22% just to the right of the cutoff.

The reason why the probability of enrollment does not jump to 100% is that students only receive offers from the highest ranked college-major in their rank-ordered list for which they are better ranked than the last admitted student. As such there may be several college-majors for which a student is better ranked than the last admitted student and yet does not receive an offer because he or she has an offer from a better ranked college-major. Note that around the admission cutoff, the *intensity* of treatment is essentially equal to one student, since panel A (probability of enrollment) and panel B (number of enrolled) largely coincide. Appendix Figure D.1 shows exactly the share of cases with one, two or more enrolled students in t at each distance to the last admitted student. Our setting only allows us to estimate spillovers for high schools going from zero enrolled students to one enrolled student in the college-major. We leave it to future research to determine how (i) the intensity of treatment, i.e., from zero to two or three enrolled students, and (ii) marginal increases above one enrolled student, i.e., from one to two or two to three enrolled students, impact cross-cohort spillovers.

Main specification. We stack thousands of college-major-year specific regression discontinuities such that for all college-major-years the rank of the last admitted student is zero. The following equation describes our baseline specification:

$$y_{s_{jt+1}} = \beta \text{distance to last admitted student}_{s_{jt}} + \gamma (\text{distance to last admitted student} \geq 0)_{s_{jt}} + \delta \text{distance to last admitted student}_{s_{jt}} \times (\text{distance to last admitted student} \geq 0)_{s_{jt}} + \mu_{jt} + \epsilon_{s_{jt+1}} \quad (2.1)$$

$y_{s_{jt+1}}$ indicates whether high school s with a marginally admitted student to college-major

¹¹In the data, we observe some cases of students accepting an offer despite being ranked below the last admitted student, though these are exceedingly rare (0.04% of our sample).

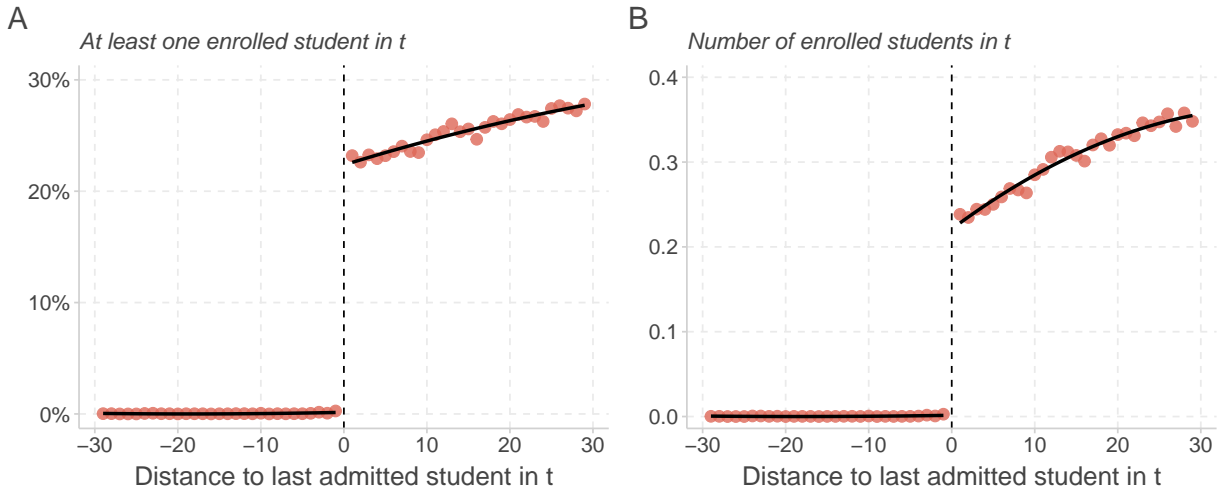


Figure 2.1. Probability of Older Schoolmate Enrolling in College-Major as a Function of the Distance to the Last Admitted Student

Notes: This figure shows the probability that a high school has at least one enrolled student (panel A) and its number of enrolled students (panel B) in the college-major as a function of the distance to the last admitted student in the same year. The distance to the last admitted student is defined as the rank of the high school's best ranked applicant by the college-major centered around the rank of the college-major's last admitted student.

j in year t has more applications and enrollments in college-major j in year $t + 1$. *distance to last admitted student* is the running variable described above, and (*distance to last admitted student* ≥ 0) is a dummy variable indicating whether an older schoolmate was ranked above the last admitted student by college-major j in year t . The interaction term allows the slopes to differ around the cutoff. μ_{jt} are college-major-year fixed effects as recommended by Fort et al. (2022). Thus the identifying variation comes from differences in *exposure* to a given college-major across high schools. ϵ_{sjt} is an error term. γ is our coefficient of interest.

This specification estimates intent to treat effects, i.e. the effect of having an older schoolmate ranked above the last admitted student to the college-major but not necessarily enrolling in it. To estimate the effect of an older schoolmate's actual enrollment, we instrument enrollment with an indicator for being ranked above the rank of the last admitted student. The 2SLS estimate corresponds to the ratio between the intent to treat estimate and the first-stage estimate. As we discussed in Section 2.2, we do not observe actual enrollment but rather whether a student *accepts* a college-major's offer on the application platform. This leaves our intent to treat estimates unchanged, however it would affect our 2SLS estimates. Since the first-stage for actual enrollment is necessarily smaller than

the first-stage for accepting an offer, if anything the 2SLS coefficients would be larger, and thus the reported instrumented estimates are lower bounds on the effect of an older schoolmate’s actual enrollment.

Following [Cattaneo et al. \(2019\)](#)’s guidelines, we estimate the coefficient of interest non-parametrically using local linear regressions. Specifically, linear regressions are fit on both sides of the threshold using a triangular kernel which gives more weight to observations near the threshold. We compute MSE-optimal bandwidths following [Calonico et al. \(2014\)](#) for our four main outcomes: (i) having at least one applicant to the same college-major as an older schoolmate, (ii) the number of applicants to this college-major, (iii) having at least one enrolled student in this college-major, and (iv) the number of enrolled students in this college-major. To ensure a constant sample across estimates, we use a common bandwidth of 21.32 ranks, which corresponds to the smallest bandwidth for these main outcomes (as done in [Altmejd et al. \(2021\)](#)). Since the same high school can potentially have a student at the margin of admission for several college-majors, we cluster standard errors at the high school - year level.

Robustness. In Appendix C, we show that our main results are robust to (i) varying the bandwidth used, and (ii) estimating discontinuities at placebo cutoffs.

3.4. *Identifying Assumptions*

The main identifying assumption is that the rank of the last admitted student to a college-major is exogenous with respect to that student’s high school. This implies two assumptions common to regression discontinuity designs: (i) college-majors cannot strategically manipulate their ranking of applicants and neither can applicants, and (ii) potential confounders do not jump discontinuously at the admission rank cutoff.

First, as discussed in Section 2.1, because college-majors do not know students’ rank ordered lists, they cannot know ex-ante which student will end up being the last admitted student.¹² That is, even though college-majors may use students’ high schools to rank them, they cannot predict from which high school the last admitted student will come from. As can be seen in Figure 2.2, we find no clear indication of manipulation of the running variable. Appendix Figure D.2 shows the composition in terms of type of program

¹²Technically, for oversubscribed public university programs, students’ rank-ordered lists were used to randomly allocate students to such programs. Nonetheless, these programs could not predict ex-ante the high school of origin of the last admitted student. See [Bechichi et al. \(2021\)](#) for additional details on this lottery system.

at each value of the running variable.

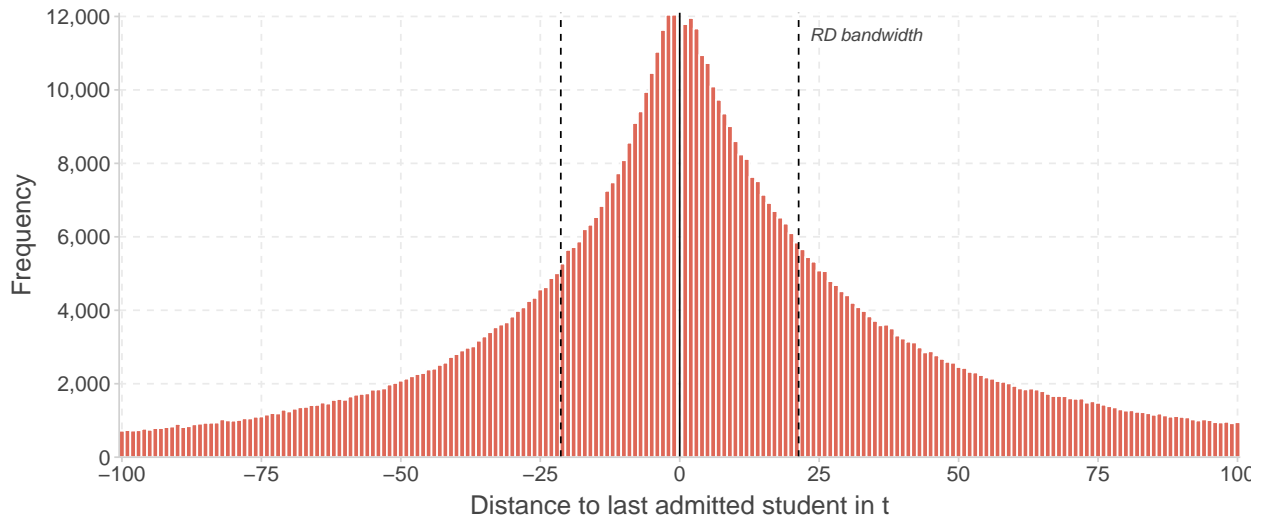


Figure 2.2. Distribution of Running Variable

Notes: This figure shows the distribution of the running variable which corresponds to the rank of the high school's best ranked applicant by the college-major centered around the rank of the college-major's last admitted student. The dashed lines represent the the regression discontinuity (RD) bandwidth used in the analysis.

Second, we test whether characteristics of high schools are identical around the rank of the last admitted student. Figure 2.3 displays the estimated discontinuities in high school characteristics around the rank of the last admitted student. All the differences are extremely small in magnitude and statistically insignificant. Crucially, there are no differences in the number of applicants in the treatment year, nor in the number of ranked applicants in the treatment year.¹³ The underlying figures in Appendix Figures D.3 clearly visually confirm the lack of any discontinuities in high school characteristics around the cutoff. In Appendix Figures D.4 we also show there are no differences in the characteristics of college-majors around the cutoff. These differences would in any case be absorbed by the college-major-year fixed effects included in equation (2.1), but this ensures that the visual evidence we present below is valid.

4. Main Results

This section presents the main results on within-high school spillovers on higher education choices. We show that, within a given high school, younger cohorts are more likely to

¹³Note that by necessity high schools on both sides of the cutoff have at least one ranked applicant in treatment year, otherwise they would not have an applicant rank for the college-major.

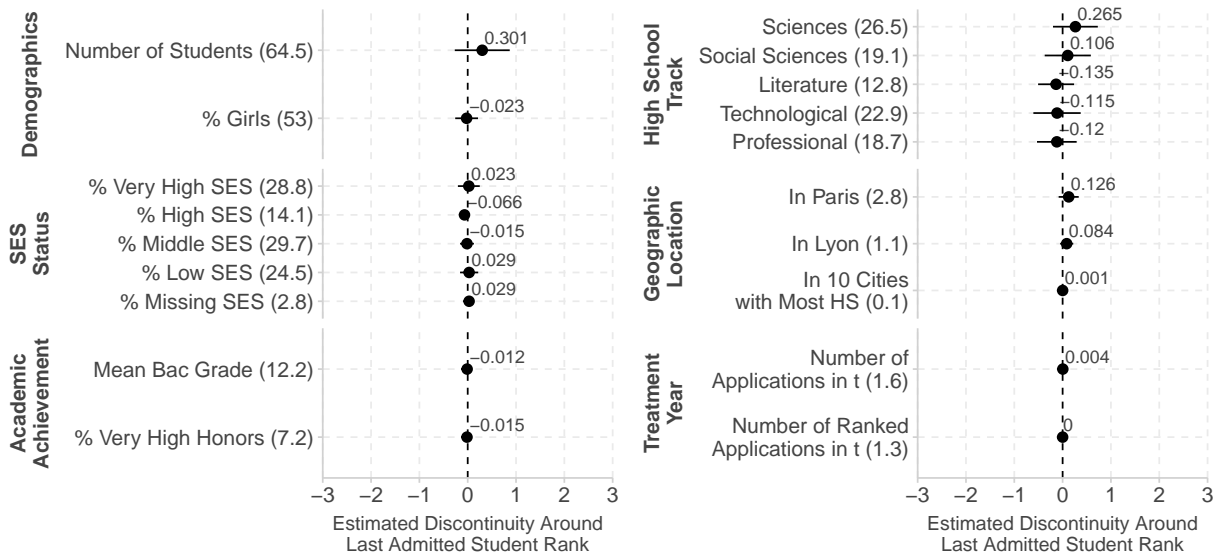


Figure 2.3. Discontinuity in High School Characteristics

Notes: This figure shows the estimates of the discontinuity in high school characteristics around the rank of college-majors' last admitted student. The high school characteristics are reported on the y-axis. The mean of the high school characteristic just below the rank of the last admitted student (e.g., $[-5, -1]$) is shown in parenthesis. All the specifications in the figure correspond to local linear regressions using a triangular kernel, and include college-major - year fixed effects. The bandwidth (21.32) corresponds to the smallest MSE-optimal bandwidth (Calonico et al., 2014) for the main college-major outcomes. Statistical significance is based on standard errors clustered at the high school - year level. Confidence intervals correspond to 95% confidence intervals. ***, **, and * indicate significance at 1, 5, and 10%, respectively.

apply to and enrol in the same college-majors as a marginally admitted older schoolmate than if the older schoolmate had been marginally rejected.

4.1. Within-High School Spillovers on Applications

Students are more likely to apply to the same college-majors as the previous cohort of students within the same high school. We illustrate this phenomenon in Figure 2.1a panels A and B. These figures show the reduced-form relationship, across all college-majors and years, between a high school's applications in $t + 1$ and the rank of this high school's best ranked student relative to the rank of the last admitted student to the college-major in year t . Panel A displays the extensive margin of applications, i.e., whether there is at least one application to the college-major, while panel B displays the intensive margin, i.e., the number of applicants to the college-major. Both figures are binned scatter plots, where each point represents the average outcome in $t + 1$ for high schools at ranks relative to the last admitted student between -30 and 30 in t .

There is a clear discontinuity in the likelihood of having at least one applicant, and in the number of applications, to college-majors for which there was a marginally admitted student from the same high school in the previous cohort. The first row of Table 2.1 reports the intent-to-treat (“Older schoolmate above cutoff (ITT)”) estimates. They suggest a high school with a marginally admitted student to a college-major in the previous cohort increases its likelihood of having at least one application to the same college-major in the following year by 1.6 percentage points, and its number of applicants by 0.046. These correspond respectively to a 5% and 6% increase relative to the mean outcome just below the cutoff. Thus the effects appear to be very slightly larger for the intensive margin of applications than for the extensive margin.

In the second row of Table 2.1, we report 2SLS estimates obtained by combining these reduced-form estimates with our first-stage results (“Older schoolmate enrolls (2SLS)”). These coefficients correspond to the effect of a student in a high school enrolling in a given college-major on the application and enrollment decisions of the following cohort of students from the same high school. Since our first-stage estimates are rather moderate, a 22 percentage point increase in the probability of treatment, the 2SLS estimates are significantly larger than the intent to treat estimates.¹⁴ A student’s enrollment in a college-major leads to a 7.3 percentage points increase in the probability of having at least one application from the same high school to the same college-major the following year, and a 0.24 increase in the number of applicants. Relative to the baseline mean, these effects are substantial. They represent a 25% and 30% increase respectively.

To get a sense of the magnitude of these effects, we can compare them with the effects found for siblings spillovers on higher education choices. [Altmejd et al. \(2021\)](#) find that an older siblings’ enrollment in a given college-major increases their younger siblings’ application to the same college-major by 36% in Chile, 32% in Croatia, and 58% in Sweden. Thus our estimated high school cohort spillovers correspond to between 43% and 78% of siblings spillovers.¹⁵ This is a sizable effect considering how much more diffuse cross-cohort spillovers are likely to be relative to one’s sibling. This highlights the essential role by students’ high school environments in shaping their higher education choices.

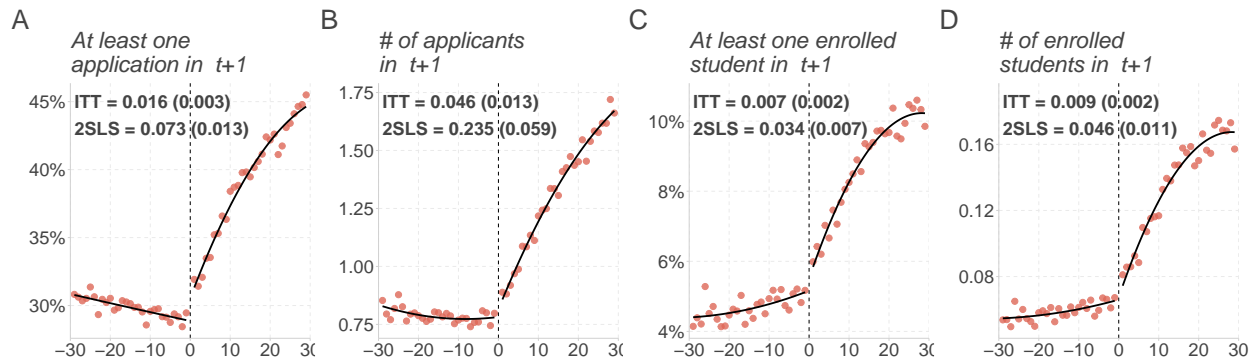
¹⁴Recall that the 2SLS estimates are defined as the ratio between the intent to treat and the first-stage estimates.

¹⁵Since no siblings spillovers have yet been estimated for France, we can only compare our within-high school spillovers with the siblings spillovers estimated in other countries.


4.2. *Within-High School Spillovers on Enrollment*

Students are also more likely to enrol in the same college-major as marginally admitted students from the same high school's previous cohort. Figure 2.1a panels C and D shows the same relationship as for applications except with enrollment as the outcome. As with applications, there is a clear discontinuity in a high school's likelihood of having at least one student enrol and in the number of enrolled students in the same college-major as a marginally admitted student from the previous cohort. Table 2.1 reports the reduced-form and 2SLS estimates of the discontinuity. Having a marginally admitted student in the previous cohort increases a high school's probability of having at least one enrollment in the college-major by 0.7 percentage points and its number of enrolled students by 0.009. These coefficients are small in absolute terms but considering how unlikely it is for a high school to have a student enrol in any given college-major, these represent 14% of the counterfactual means. The 2SLS estimates are larger by construction, yielding a 3.4 percentage point increase in the probability of having at least one enrolled student and a 0.046 increase in the number of enrolled students. Relative to the baseline mean, these correspond respectively to 67% and 72% increases.


Compared to sibling spillovers, the within-high school spillover effects on enrollment are about one third greater than in Chile (50%) and between 40% and 88% of those found in Sweden (167%) and Croatia (76%) respectively. There are two possible explanations as to why the cohort spillovers on the enrollment margin are closer to the siblings spillover estimates compared to the application margin. First, when ranking applicants, college-majors use all available information, including the applicant's high school of origin. Thus, if a college-major enrolls a student from a high school from which it had previously never accepted any applicant, it may learn about the high school's quality and therefore change how it ranks future applicants from this high school. Second, younger cohorts may rank the college-major higher up in their rank-ordered list, increasing the likelihood that they enrol should they be ranked better than the last admitted student. We plan to investigate both margins of response in a future version of the paper.



(a) College-Major Spillovers



(b) College Spillovers



(c) Major Spillovers

Figure 2.1. Within-High School Spillovers on Applications to and Enrollment in College-Major, College and Major with Marginally Admitted Older Schoolmate

Notes: This figure shows non-parametric binned scatter plots of the relationship between high schools' application and enrollment outcomes for a college-major in $t + 1$ and high schools' distance to the college-majors' last admitted student in t . The application and enrollment outcomes are reported in the figure facet title. Each point corresponds to the average outcome value for high schools with distance to the last admitted student equal to the value on the x-axis. The fitted lines correspond to second-order polynomial fits through the conditional expectation. For each figure, we report the intent-to-treat (ITT) and instrumented (2SLS) estimates, which include college-major - year fixed effects. All the specifications in the figure correspond to local linear regressions using a triangular kernel, and include college-major - year fixed effects. The bandwidth (21.32) corresponds to the smallest MSE-optimal bandwidth (Calonico et al., 2014) for the main college-major outcomes. Standard errors clustered at the high school - year level are reported in parentheses.

Table 2.1: High School Cohort Spillovers on Applications to and Enrollment in Degree, College and Major with Marginally Admitted Student

	College-Major Spillovers						College Spillovers						Major Spillovers					
	Applications		Enrollment		Applications		Enrollment		Applications		Enrollment		Applications		Enrollment			
	At least one (1)	Number (2)	At least one (3)	Number (4)	At least one (5)	Number (6)	At least one (7)	Number (8)	At least one (9)	Number (10)	At least one (11)	Number (12)						
Older schoolmate above cutoff (ITT)	0.016*** (0.003)	0.046*** (0.013)	0.007*** (0.002)	0.009*** (0.002)	0.017*** (0.003)	0.563*** (0.068)	0.02*** (0.003)	0.171*** (0.023)	0.003 (0.003)	-0.005 (0.046)	0	-0.015 (0.014)						
% of counterfactual mean	5.37	5.88	14.07	14.08	3.13	11.99	9.38	19.27	0.4	-0.1	0.11	-1.36						
Older schoolmate enrolls (2SLS)	0.073*** (0.013)	0.235*** (0.059)	0.034*** (0.007)	0.046*** (0.011)	0.096*** (0.017)	3.158*** (0.365)	0.109*** (0.014)	0.96*** (0.125)	0.03 (0.02)	0.214 (0.364)	0.016 (0.022)	-0.065 (0.11)						
% of counterfactual mean	25.16	30.07	67.23	72.05	17.37	67.23	52.35	108.16	4.73	4.47	4.65	-6.1						
College-major-year FE	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓						
Obs. (right)	180,974	180,974	180,974	180,974	180,974	180,974	180,974	180,974	180,974	180,974	180,974	180,974						
Obs. (left)	173,194	173,194	173,194	173,194	173,194	173,194	173,194	173,194	173,194	173,194	173,194	173,194						
Counterfactual mean [-5,-1]	0.29	0.782	0.05	0.064	0.552	4.698	0.208	0.887	0.637	4.781	0.354	1.069						
Bandwidth	21.32	21.32	21.32	21.32	21.32	21.32	21.32	21.32	21.32	21.32	21.32	21.32						
First stage	0.224*** (0.002)	0.224*** (0.002)	0.224*** (0.002)	0.224*** (0.002)	0.182*** (0.003)	0.182*** (0.003)	0.182*** (0.003)	0.182*** (0.003)	0.126*** (0.003)	0.126*** (0.003)	0.126*** (0.003)	0.126*** (0.003)						
First stage F-stat	18,863	18,863	18,863	18,863	8,016	8,016	8,016	8,016	2,945	2,945	2,945	2,945						

Notes: This table reports estimates of within-high school spillovers. "At least one application" refers to the probability that a high school has at least one application to the college-major, college or major in $t + 1$ of a marginally admitted older schoolmate in t ; "number of applications" refers to the number of applications to the marginally admitted students' college-major, college or major. The same applies for the enrollment outcomes. The first row for each outcome presents intent-to-treat estimates, while the second row presents 2SLS estimates in which older schoolmates' enrollment is instrumented with them being ranked above the admission rank cutoff. All the specifications in the table correspond to local linear regressions using a triangular kernel, and include college-major - year fixed effects. The bandwidth (21.32) corresponds to the smallest MSE-optimal bandwidth (Calónico et al., 2014) for college-major outcomes. Standard errors clustered at the high school - year level are reported in parentheses. ***, **, * and * indicate significance at 1, 5, and 10%, respectively.

4.3. Cross Cohort Spillovers on College and Major

Students' enrollment decisions could also have broader spillovers on the college or major applications and enrollments of students in the subsequent cohort of the same high school. So far we have shown there are sizable college-major-specific spillovers from one high school cohort to the next. Each college-major corresponds to a college and a major and thus either component could influence subsequent students' higher education decisions.

We undertake the exact same analysis as for college-majors, this time for colleges, and for majors separately. Due to the specificity of the French higher education system, majors are partly college-specific because some tracks are only offered in some types of institutions. For example, technical college-majors are only offered in University Institutes of Technology (*Instituts Universitaires de Technologie (IUT)*) and vocational college-majors are only offered in Sections of Superior Technicians (*Sections de Techniciens Supérieurs (STS)*).

Figures 2.1b and 2.1c and Table 2.1 report the results. We find sizable within-high school spillovers on applications and enrollments in a given institution but no spillovers on majors. The higher education institution-specific spillovers are particularly pronounced for the intensive margin of applications and enrollments: a 0.56 increase in the number of applicants and 0.17 increase in the number of enrolled students. Relative to the counterfactual mean, these are equivalent to a 12% and 19% increase. Since the first-stage is moderate (0.18) the 2SLS estimates are consequently even larger.

As with within-high school spillovers in college-majors, we can get a sense of the magnitude of these estimates by comparing them to sibling spillovers. On the application margin, our within-high school college spillovers correspond to between 27% (Sweden) and 102% (Chile) of the siblings effect, and 34% (Sweden) to 182% (Chile) on the enrollment margin.

4.4. Robustness

Bandwidth size. Appendix Figure C.1 shows our baseline estimates are insensitive to the choice of bandwidth. Specifically, the estimates for spillovers in college-majors are very stable, while college-specific spillovers are slightly increasing with the bandwidth. The estimates for major spillovers are consistently very small and insignificant.

Placebo cutoffs. Appendix Figure C.2 shows our baseline estimates are robust to placebo

thresholds. Specifically we estimate discontinuities at other locations along the running variable and find the only significant effects are found at the actual cutoff.

Placebo outcomes. We estimate discontinuities in our four main outcomes measured in the year prior to the treatment. We should observe no effects for these outcomes because by construction they happen before the treatment takes place. Figure 2.2 shows the results. We focus on treatment years 2014-2015 because these are the only years for which we can measure outcomes in $t-1$, t (treatment year), $t+1$ (baseline outcomes) and $t+2$ for a constant sample of high schools. The coefficients for the extensive margin (at least one application or admitted student) are small though significant which is unexpected. It is not clear how such a discontinuity could arise, and in fact the graphical evidence presented in Appendix Figures D.5 shows no clear indication of discontinuities for pre-treatment year for any outcome. Nonetheless, the estimates remain smaller than our baseline coefficients. Regarding the intensive margin (number of applicants and admitted students) the coefficients are very small and insignificant.

College-majors with no applicants in t-1. To further strengthen the validity of our approach, we re-estimate our main results on the subsample of college-majors for which there were no applicants in the $t - 1$ cohort of students, i.e., in the cohort just before the treatment cohort. As such we are certain there can be no issue with the previous placebo analysis in this setting because there are no applicants in the placebo year. The results of this analysis are shown in Table 2.2, with the underlying figures displayed in Appendix Figure D.6. The estimates are all significant and of slightly larger magnitude than the baseline results. One explanation for why these effects are larger than the baseline ones could be that precisely because the focus here is only on college-majors for which the high school had no applicants in $t - 2$, having a marginally student may disproportionately increase the salience of this college-major.

4.5. *Snowball Effect*

So far, we have found that students' enrollment decisions affect the higher education choices of students in the following cohort of the same high school. It's possible that these one time events have long-lasting effects on the cohort of students two years later or even three years later. As shown in Figure 2.2 we find evidence that the within-high school spillovers persist over time. The coefficients in $t + 2$, i.e., two years after treatment, are a bit smaller than the coefficients in $t + 1$ but still sizable and statistically significant. For the extensive margin of applications, the $t + 2$ spillovers are 94% as large those in $t + 1$,

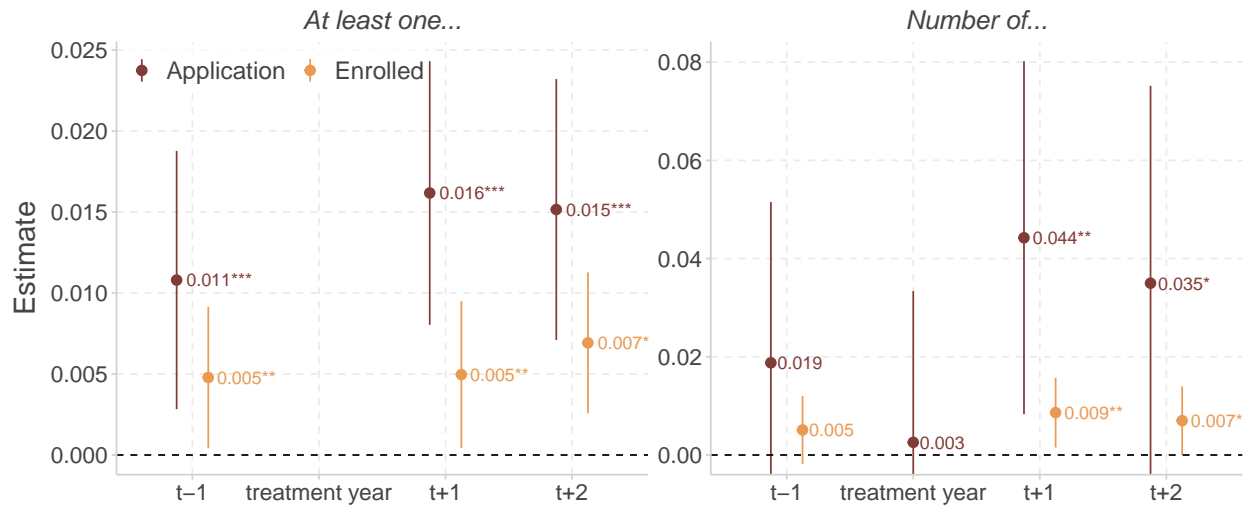


Figure 2.2. Placebo and Snowball Effects

Notes: This figure shows estimates of within-high school spillovers in applications and enrollments for outcomes measured in different years. “ $t - 1$ ” refers to the year prior to the older schoolmate’s marginal admission, “treatment year” refers to the year of an older schoolmate’s marginal admission, and $t + 1$ and $t + 2$ correspond, respectively to high schools’ application and enrollment outcomes one and two years following the marginal admission of one of its students. The sample is restricted to treatment years 2014 and 2015 to ensure the sample is constant across estimates. All the specifications in the figure correspond to local linear regressions using a triangular kernel, and include college-major - year fixed effects. The bandwidth (21.32) corresponds to the smallest MSE-optimal bandwidth (Calonico et al., 2014) for the main college-major outcomes. Statistical significance is based on standard errors clustered at the high school - year level. Confidence intervals correspond to 95% confidence intervals. ***, **, and * indicate significance at 1, 5, and 10%, respectively.

though there are few applicants. This result is important because it highlights that higher education choices within high schools are quite path dependent and that seemingly small shocks can continue to influence students at least two years after the initial “shock”.

4.6. Heterogeneity

Before trying to tease out the mechanisms underlying within-high school spillovers, we explore the heterogeneity of the effects across a number of dimensions. First, we find that the estimates are surprisingly stable over the four year used in the analysis, underscoring that the spillover we uncover are not restricted to the idiosyncracies of a given year. This is suggestive that the within-high school spillovers we have identified are a structural determinant of students’ higher education choices. Next, we assess the extent to which spillovers vary by students’ characteristics, high schools’ characteristics, college-majors’ characteristics, and the interaction between high schools’ and college-majors’ characteristics. We discuss each of these in turn below.

Table 2.2: College-Major Spillovers on Applications to and Enrollment for College-Majors with No Applications from High School in $t-1$

	Applications		Enrollment	
	At least one (1)	Number (2)	At least one (3)	Number (4)
Older schoolmate above cutoff (ITT)	0.012*** (0.004)	0.031*** (0.009)	0.004*** (0.002)	0.005** (0.002)
% of counterfactual mean	6.68	9.76	17.27	18.53
Older schoolmate enrolls (2SLS)	0.061*** (0.018)	0.156*** (0.045)	0.02*** (0.008)	0.025** (0.01)
% of counterfactual mean	33.96	49.82	87.85	94.26
College-major-year FE	✓	✓	✓	✓
Obs. (right)	85,271	85,271	85,271	85,271
Obs. (left)	89,121	89,121	89,121	89,121
Counterfactual mean	0.179	0.314	0.023	0.027
Bandwidth	21.32	21.32	21.32	21.32
First stage	0.198*** (0.003)	0.198*** (0.003)	0.198*** (0.003)	0.198*** (0.003)
First stage F-stat	7,477	7,477	7,477	7,477

Notes: This table reports estimates of within-high school spillovers for the subset of college-majors for which there were no applicants in a high school's $t - 1$ cohort of students. The first row for each outcome presents intent-to-treat estimates, while the second row presents 2SLS estimates in which older schoolmates' enrollment is instrumented with them being ranked above the admission rank cutoff. All the specifications in the table correspond to local linear regressions using a triangular kernel, and include college-major - year fixed effects. The bandwidth (21.32) corresponds to the smallest MSE-optimal bandwidth (Calonico et al., 2014) for college-major outcomes. Standard errors clustered at the high school - year level are reported in parentheses. ***, **, and * indicate significance at 1, 5, and 10%, respectively.

Students' characteristics. We start by exploring how spillovers vary depending on students' characteristics. Thus, here we are interested in understanding whether some types of students are more susceptible to follow a marginally admitted older schoolmate. The results are presented in Figure 2.3. There are no differences in spillovers on applications between boys and girls, but the enrollment effects are larger for boys. This could be due to the way boys rank their applications or to the way college-majors respond to their applications. Additionally, low socioeconomic status students are only slightly more responsive than their very high SES peers. This is surprising as one may suspect that low SES students would exhibit greater responses since they are presumably less informed

about higher education programs and therefore may be more susceptible to follow an older schoolmate’s higher education trajectory.




Figure 2.3. Heterogeneity by Student Characteristics

Notes: This figure shows estimates of within-high school spillovers in applications and enrollments for subsamples of students. The application and enrollment outcomes are reported in the figure facet titles, while the subsamples are reported on the y-axis. Socioeconomic status (SES) is based on students’ legal guardian’s occupation. All the specifications in the figure correspond to local linear regressions using a triangular kernel, and include college-major - year fixed effects. The bandwidth (21.32) corresponds to the smallest MSE-optimal bandwidth (Calonico et al., 2014) for the main college-major outcomes. Statistical significance is based on standard errors clustered at the high school - year level. Confidence intervals correspond to 95% confidence intervals. ***, **, and * indicate significance at 1, 5, and 10%, respectively. The percentage change relative to the counterfactual mean ($[-5, -1]$) is shown in parenthesis next to the estimates.

High schools’ characteristics. The results are displayed in Figure 2.4. All high school tracks display spillovers of roughly equal magnitude, with slightly larger ones (in percentage terms) for the literature track. It is interesting that all high school tracks exhibit cross-cohort spillovers as this suggests that information acquisition about higher education programs is relevant in varied contexts. That being said, within-high school spillovers are largest in small high schools (between 0 and 30 students), and are decreasing in high school size. This could be the result of several non-mutually exclusive mechanisms. Small high schools may exhibit closer relationships between students and




Figure 2.4. Heterogeneity by High School Characteristics

Notes: This figure shows estimates of within-high school spillovers in applications and enrollments for subsamples of high schools. The application and enrollment outcomes are reported in the figure facet titles, while the subsamples are reported on the y-axis. High schools' academic level is defined as the median of its students' end of high school exam (Bac) grades. The quintiles of academic level are calculated over *all* high schools in the full sample, not only among high schools in the regression discontinuity sample. All the specifications in the figure correspond to local linear regressions using a triangular kernel, and include college-major - year fixed effects. The bandwidth (21.32) corresponds to the smallest MSE-optimal bandwidth (Calonico et al., 2014) for the main college-major outcomes. Statistical significance is based on standard errors clustered at the high school - year level. Confidence intervals correspond to 95% confidence intervals. ***, **, and * indicate significance at 1, 5, and 10%, respectively. The percentage change relative to the counterfactual mean ($[-5, -1]$) is shown in parenthesis next to the estimates.

their teachers, and thus teachers might be better aware of their students' higher education choices. As such they may encourage their subsequent classes to apply to the same college-majors as their past students. Another explanation could be that smaller high schools maintain better links with their alumni through, for example, annual alumni gatherings/forums. A last explanation could simply be that smaller high schools are located in more rural areas where information about college-majors may be more scarce and therefore informational shocks are amplified to a much larger extent than in information-

rich high schools located in large cities.

We also explore whether high schools' academic level matters for the intensity of spillovers. We measure high school's academic level as the median of its students' end of high school exam grades, and then group high schools into quintiles. The quintiles of academic level are calculated over *all* high schools in the full sample, not only among high schools in the regression discontinuity sample. Our results does not suggest a very obvious story: high schools in the second quintile and in the top 20% display spillovers on the extensive margin of applications, but only high schools in quintiles 2 and 3 have any spillovers on the intensive margin of applications. To some extent, this suggests that high schools' academic level in itself does not appear to predict within-high school spillovers.

College-majors' characteristics. Third, we examine whether some college-majors induce larger within-high school spillovers. The largest spillovers are found for public university, technical and professional programs, and there are no spillovers for preparatory classes, which tend to be quite prestigious, or other college-majors. In line with this result, only college-programs in the lower part of the selectivity distribution (proxied using the median end-of-high school exam grade of enrolled students¹⁶) exhibit spillovers on applications. This is intriguing because one could have expected very selective college-majors to induce the largest spillovers and yet we found quite the opposite. This tentatively suggests to us that informational barriers may be more prevalent in our setting than aspirational barriers.

Interaction between high schools' and college-majors' characteristics. Lastly, we assess how the interaction between high schools' and college-majors' characteristics may shape within-high school spillovers. The results are shown in Figure 2.6. First, we estimate how the geographic distance¹⁷ between the high school and the college-major might affect the likelihood of students following the marginally admitted older schoolmate. It is not obvious how distance would affect spillovers because closer college-majors are more likely to be known by students though if they are not they do not imply additional living costs. Conversely, far away college-majors are less likely to be known by students but they would imply incurring potentially important living costs. The results point towards

¹⁶The deciles of selectivity are calculated over *all* college-majors in the full sample, not only among college-majors in the regression discontinuity sample.

¹⁷We obtain high schools' and college-majors' precise geographic location (longitude and latitude) from available open data. We are currently missing exact geographic location for 7% of high schools and for 4% of college-majors.




Figure 2.5. Heterogeneity by College-Major Characteristics

Notes: This figure shows estimates of within-high school spillovers in applications and enrollments for subsamples of college-majors. The application and enrollment outcomes are reported in the figure facet titles, while the subsamples are reported on the y-axis. College-majors' selectivity is measured as median end-of-high school exam grade of enrolled students. The deciles of selectivity are calculated over *all* college-majors in the full sample, not only among college-majors in the regression discontinuity sample. All the specifications in the figure correspond to local linear regressions using a triangular kernel, and include college-major - year fixed effects. The bandwidth (21.32) corresponds to the smallest MSE-optimal bandwidth (Calonico et al., 2014) for the main college-major outcomes. Statistical significance is based on standard errors clustered at the high school - year level. Confidence intervals correspond to 95% confidence intervals. ***, **, and * indicate significance at 1, 5, and 10%, respectively. The percentage change relative to the counterfactual mean ($[-5, -1]$) is shown in parenthesis next to the estimates.

geographically close (less than 25 km) and moderately far (between 50 and 100 km) programs induce the largest spillovers. In terms of the intensive margin of applications and enrollment on these moderately far programs display significant cross-cohort spillovers.

Additionally, we estimate whether the difference between the academic level of the high school and the college major might impact spillovers. Specifically, we use the measures of high school academic level (proxied using its students' median end-of-high school exam grade) and college-major selectivity (proxied using the median end-of-high school exam grade of enrolled students) defined previously, and restrict the analysis to high schools and college-majors in the bottom and top quartile of their respective academic level distributions. We then estimate spillovers for every combination across the two groups




Figure 2.6. Heterogeneity by High School *and* College-Major Characteristics

Notes: This figure shows estimates of within-high school spillovers in applications and enrollments for subsamples of high schools and college-majors. The application and enrollment outcomes are reported in the figure facet titles, while the subsamples are reported on the y-axis. See Figures 2.4 and 2.5's notes for details on the definitions of high schools academic level and college-major selectivity, used for the *Diff. Academic Level* results. All the specifications in the figure correspond to local linear regressions using a triangular kernel, and include college-major - year fixed effects. The bandwidth (21.32) corresponds to the smallest MSE-optimal bandwidth (Calonico et al., 2014) for the main college-major outcomes. Statistically significance is based on standard errors clustered at the high school - year level. Confidence intervals correspond to 95% confidence intervals. ***, **, and * indicate significance at 1, 5, and 10%, respectively. The percentage change relative to the counterfactual mean ($[-5, -1]$) is shown in parenthesis next to the estimates.

of high schools (Top 25% and Bottom 25%) and the two groups of college-majors (Top 25% and Bottom 25%). Somewhat intuitively we find that students in the bottom 25% of high schools' academic level are significantly more responsive to an older schoolmate marginally admitted to a college-major in the top 25% of selectivity. This effect could be interpreted as raising the aspirations or awareness of these high schools' top performing students. Conversely, we find intriguingly that spillovers are quite large for top 25% high schools for whom the marginally admitted student went to a college-major in the bottom quartile of selectivity. It is difficult to understand this result without knowing more about which students are induced to apply.

5. Mechanisms

We study two potential non-mutually exclusive mechanisms that could underpin the within-high school spillovers we uncovered: (i) the role of teachers, and (ii) student role model effects. Teachers may be aware of their past students' higher education choices and thus recommend them to their current students. Moreover, students may be more likely to follow in the steps of older schoolmates who share similar characteristics such as gender or socioeconomic background. This latter mechanism to some extent helps understand what type of policy intervention would most resemble this spillover.

5.1. *Role of Teachers*

Because we do not observe all of students' teachers, we assess the role teachers may play in two indirect ways. First, we exploit the fact that, in France, each senior class is assigned a "principal" teacher. These principal teachers, in addition to teaching their subject-area, are in charge of all administrative duties for the class. In particular, as part of these duties, they are in charge of assisting students with their higher education applications. As such, these teachers tend to be well-aware of their students' higher education choices relative to their class' non-principal teachers. Therefore, we examine whether students who share the principal teacher as the marginally admitted older schoolmate are more likely to apply to the older schoolmate's college-major than students who do not share the same principal teacher. If teachers are the mediating factor, we would expect same-principal teacher students to exhibit larger spillovers than non-same-principal teacher students.

Second, we use students' class identifier number, e.g., senior class A, senior class B, etc. as a proxy for the set of teachers they are likely to have each year. The assumption here is that senior class A's teachers will be the same from one cohort to the next. This is admittedly a highly imperfect proxy for the set of teachers but we think in practice there is quite likely to be some overlap from one year to the next. As such, we compare spillovers for students sharing the same class number as the enrolled older schoolmate with students in a different class number. For both of these teacher analyses, we focus on the subsample of high schools with at least two classes (63% of high schools).

Figure 2.1 displays the results of this analysis. The estimates for students sharing the same principal teacher or class number as the marginally admitted older schoolmate are essentially the same as for students not sharing these teachers. We interpret these results as suggestive evidence that teachers do not appear to play a role in explaining our within-high school spillovers. This could be explained by the fact that students do not

necessarily seek out information about higher education programs from their teachers and perhaps that teachers are more likely to recommend a large variety of college-majors rather than only the subset of programs attended by their past students.




Figure 2.1. Older Schoolmate Teacher-Specific Spillovers

Notes: This figure shows estimates of within-high school spillovers in applications for students sharing the same *principal* teacher and class number as the marginally admitted older schoolmate and for students who do not. The application outcomes are reported in the figure facet titles. All the specifications in the figure correspond to local linear regressions using a triangular kernel, and include college-major - year fixed effects. The bandwidth (21.32) corresponds to the smallest MSE-optimal bandwidth (Calonico et al., 2014) for the main college-major outcomes. Statistical significance is based on standard errors clustered at the high school - year level. Confidence intervals correspond to 95% confidence intervals. ***, **, and * indicate significance at 1, 5, and 10%, respectively. The percentage change relative to the counterfactual mean ($[-5, -1]$) is shown in parenthesis next to the estimates.

5.2. Information vs Role Model

Next, we examine the role played by role models in shaping within-high school spillovers. Indeed, students may not be as responsive depending on the characteristics of the marginally admitted student from the previous graduating cohort. The literature on role models has shown that female professors have large impacts on girls' performance in maths/science and likelihood to pursue a STEM degree (Carrell et al., 2010), as well as graduating with an economics degree (Canaan and Mouganie, 2021). The same types of effects have been found for Black students randomly assigned to a Black teacher (Gershenson et al., 2022). Girls are also induced to apply to STEM degrees by external female scientists (Breda et al., 2023).

In spirit of these studies, we estimate whether girls are more likely to follow a marginally

admitted older schoolmate if she was a girl rather than a boy, and if boys are more likely to follow boys. Moreover, though the literature on socioeconomic status role models are less common, we estimate same-SES effects to assess whether low SES students are more likely to follow a low SES than a high SES older schoolmate. This analysis helps us better understand the mechanisms that could explain the within-high schools spillover effects presented above. Since all students within a high school are likely to have the same information about the admission outcomes of students from the previous cohort, evidence for role model effects would suggest that older schoolmates serve as role models for their younger high school peers. The results from this analysis are presented in Figure 2.2, with the underlying visual evidence in Appendix Figures D.7-D.10.




Figure 2.2. Older Schoolmate Gender- and SES-Specific Spillovers

Notes: This figure shows estimates of within-high school spillovers in applications for students sharing the same gender and SES as the marginally admitted older schoolmate and for students who do not. The application outcomes are reported in the figure facet titles, while the marginally admitted older schoolmate’s characteristic are reported on the y-axis. Socioeconomic status (SES) is based on students’ legal guardian’s occupation. All the specifications in the figure correspond to local linear regressions using a triangular kernel, and include college-major - year fixed effects. The bandwidth (21.32) corresponds to the smallest MSE-optimal bandwidth (Calonico et al., 2014) for the main college-major outcomes. Statistical significance is based on standard errors clustered at the high school - year level. Confidence intervals correspond to 95% confidence intervals. ***, **, and * indicate significance at 1, 5, and 10%, respectively. The percentage change relative to the counterfactual mean $([-5, -1])$ is shown in parenthesis next to the estimates.

Gender role models. Students sharing the same gender as the marginally admitted older schoolmate are significantly more likely to follow him or her. For example, girls are

1.7 percentage points (+9%) more likely to apply to the same college-major as a marginally admitted girl but are not more likely for a boy. The exact opposite results are found for boys: they are 1.4 percentage points (+8%) more likely to apply to a college-major of an older schoolmate who is a boy, but do not follow girls. These effects are also found for the intensive margin (number) of applications.

SES role models. We also find evidence of SES role model effects, except in this case only low SES students follow a low SES student while very high SES students are unresponsive regardless of whether the marginally admitted older schoolmate was low SES or very high SES. Low SES students are 1.7 percentage points (+13%) more likely to follow a marginally admitted low SES student but are unresponsive for very high SES students (0.2 percentage points). The fact that very high SES students are unresponsive regardless of the SES of the marginally admitted older schoolmate is consistent with very high SES students being better informed about higher education programs.

6. Conclusion

This paper presents the first causal evidence for within-high school spillovers in college-major choice. Specifically we find that students are likely to apply to and enrol in a college-major if a student from the same high school was admitted to this exact same college-major the previous year. We identify these spillovers by exploiting admission cutoffs generated by the French centralised application and admission system, and implementing a fuzzy regression discontinuity design using administrative data from the application platform for 2013-2017.

Specifically, students are about 1.6 percentage points (5%) more likely to apply to the same college-major as an older schoolmate, and 0.7 percentage points (14%) more likely to enrol. These estimated spillovers are very large, roughly 60% of those found for sibling spillovers in college-major, highlighting the essential role played by the high school environment in shaping students' higher education choices. In particular, these effects appear to be driven by student role model effects rather than by the role of teachers. Indeed, girls are significantly more likely to follow marginally admitted girls from the previous graduating cohort, but don't follow boys. The exact same of behavior is observed for boys. Interestingly, we also find that low SES students are more likely to follow a low SES older schoolmate, but not a very high SES older schoolmate.

We also find that the spillovers persist over time and continue in the graduating cohort

two years following marginal admission. This underscores how seemingly small shocks can snowball into persistent patterns.

For the moment, this paper leaves a very important question unanswered: do these spillovers matter? Perhaps, the spillovers we identify are simply substitutions from one college-major to a very similar other college-major, in which case, from a policy or welfare perspective it is not clear whether they matter. There are two other scenarios. First, these spillovers enable students to improve their higher education outcomes by either being steered towards actually enrolling in higher education or towards more selective/prestigious higher education programs. Second, and more pessimistically, students may follow older schoolmates even if their college-majors represents a “worse” program (whatever definition of “worse” is adopted).

We plan on addressing this question in a future version of the paper. Specifically we want to start by answering an easy question: does marginal admission of an older schoolmate causally influence the selectivity of the college-majors to which younger cohorts enrol in. This can be implemented easily and would shed some light on the previous question. Another way to answer the question is to re-run the allocation algorithm and changing treated high schools’ applications such that they do not include the college-major with a marginally admitted older schoolmate. In this counterfactual, which college-majors would students be offered a seat in? How does that college-major compare with the college-major where they actually enrolled? These questions are fundamental and we plan on tackling them as soon as possible.

References

- Aguirre, Josefa and Juan Matta**, “Walking in Your Footsteps: Sibling Spillovers in Higher Education Choices,” *Economics of Education Review*, February 2021, 80, 102062.
- Altmejd, Adam**, “Inheritance of Fields of Study,” 2023.
- , **Andrés Barrios-Fernández, Marin Drlje, Joshua Goodman, Michael Hurwitz, Dejan Kovac, Christine Mulhern, Christopher Neilson, and Jonathan Smith**, “O Brother, Where Start Thou? Sibling Spillovers on College and Major Choice in Four Countries,” *The Quarterly Journal of Economics*, June 2021, 136 (3), 1831–1886.
- Altonji, Joseph G., Erica Blom, and Costas Meghir**, “Heterogeneity in Human Capital Investments: High School Curriculum, College Major, and Careers,” *Annual Review of Economics*, September 2012, 4 (1), 185–223.
- Aucejo, Esteban M., Claudia Hupkau, and Jenifer Ruiz-Valenzuela**, “Where versus What: College Value-Added and Returns to Field of Study in Further Education,” *Journal of Human Resources*, October 2022, pp. 0620–10978R1.
- Avdeev, Stanislav, Nadine Ketel, Hessel Oosterbeek, and Bas van der Klaauw**, “Spillovers in Fields of Study: Siblings, Cousins, and Neighbors,” Technical Report 2023.
- Barrios-Fernández, Andrés**, “Neighbors’ Effects on University Enrollment,” *American Economic Journal: Applied Economics*, July 2022, 14 (3), 30–60.
- Barrios-Fernandez, Andres**, “Peer Effects in Education,” 2023.
- Bechichi, Nagui and Georgia Thebault**, “Students’ Preferences, Capacity Constraints and Post-Secondary Achievements in a Non-Selective System,” Technical Report N° G2021/01, INSEE, Paris January 2021.
- , **Julien Grenet, and Georgia Thebault**, “D’Admission Post-Bac à Parcoursup : Quels Effets Sur La Répartition Des Néo-Bacheliers Dans Les Formations d’enseignement Supérieur ?,” Insee Références Édition 2021, INSEE, Paris 2021.
- Breda, Thomas, Julien Grenet, Marion Monnet, and Clémentine Van Effenterre**, “How Effective Are Female Role Models in Steering Girls Towards STEM? Evidence from French High Schools,” *The Economic Journal*, June 2023, 133 (653), 1773–1809.
- Calonico, Sebastian, Matias D. Cattaneo, and Rocio Titiunik**, “Robust Nonparametric Confidence Intervals for Regression-Discontinuity Designs: Robust Nonparametric Confidence Intervals,” *Econometrica*, November 2014, 82 (6), 2295–2326.
- Canaan, Serena and Pierre Mouganie**, “Does Advisor Gender Affect Women’s Persistence in Economics?,” *AEA Papers and Proceedings*, May 2021, 111, 112–116.

- Carrell, Scott and Bruce Sacerdote**, "Why Do College-Going Interventions Work?," *American Economic Journal: Applied Economics*, July 2017, 9 (3), 124–151.
- Carrell, Scott E, Marianne E Page, and James E West**, "Sex and Science: How Professor Gender Perpetuates the Gender Gap," *Quarterly Journal of Economics*, 2010, 125 (3), 1101–1144.
- Cattaneo, Matias D., Nicolas Idrobo, and Rocio Titiunik**, *A Practical Introduction to Regression Discontinuity Designs: Foundations* November 2019.
- Charousset, Pauline, Fack Gabrielle, Julien Grenet, and He YingHua**, "How Do Elite Post-Secondary Programs Select Their Students?," 2021.
- Chetty, R., J. N. Friedman, N. Hilger, E. Saez, D. W. Schanzenbach, and D. Yagan**, "How Does Your Kindergarten Classroom Affect Your Earnings? Evidence from Project Star," *The Quarterly Journal of Economics*, November 2011, 126 (4), 1593–1660.
- Chetty, Raj, John N. Friedman, and Jonah E. Rockoff**, "Measuring the Impacts of Teachers II: Teacher Value-Added and Student Outcomes in Adulthood," *American Economic Review*, September 2014, 104 (9), 2633–2679.
- Estrada, Ricardo, Jeremie Gignoux, and Agustina Hatrick**, "Learning About Opportunity: Spillovers of Elite School Admissions in Peru," 2022.
- Fack, Gabrielle and Julien Grenet**, "Improving College Access and Success for Low-Income Students: Evidence from a Large Need-Based Grant Program," *American Economic Journal: Applied Economics*, April 2015, 7 (2), 1–34.
- Fort, Margherita, Andrea Ichino, Enrico Rettore, and Giulio Zanella**, "Multi-Cutoff Rd Designs with Observations Located at Each Cutoff: Problems and Solutions," *SSRN Electronic Journal*, 2022.
- Gale, D. and L. S. Shapley**, "College Admissions and the Stability of Marriage," *The American Mathematical Monthly*, January 1962, 69 (1), 9.
- Gershenson, Seth, Cassandra M. D. Hart, Joshua Hyman, Constance A. Lindsay, and Nicholas W. Papageorge**, "The Long-Run Impacts of Same-Race Teachers," *American Economic Journal: Economic Policy*, November 2022, 14 (4), 300–342.
- Hastings, Justine, Christopher Neilson, and Seth Zimmerman**, "Are Some Degrees Worth More than Others? Evidence from College Admission Cutoffs in Chile," Technical Report w19241, National Bureau of Economic Research, Cambridge, MA July 2013.
- Hoxby, Caroline and Sarah Turner**, "Expanding College Opportunities for High-Achieving, Low Income Students," SIEPR Discussion Paper 12-014, Stanford Institute for Economic Policy Research, Stanford, CA 2013.
- Hoxby, Caroline M and Sarah Turner**, "Informing Students about Their College Options: A Proposal for Broadening the Expanding College Opportunities Project," Technical Report 2013-03, The Hamilton Project 2013.

- Jackson, C Kirabo**, "What Do Test Scores Miss? The Importance of Teacher Effects on Non-Test Score Outcomes," *Journal of Political Economy*, 2018, 126 (5).
- Kirkeboen, Lars J., Edwin Leuven, and Magne Mogstad**, "Field of Study, Earnings, and Self-Selection," *The Quarterly Journal of Economics*, August 2016, 131 (3), 1057–1111.
- Kramarz, F. and O. N. Skans**, "When Strong Ties Are Strong: Networks and Youth Labour Market Entry," *The Review of Economic Studies*, July 2014, 81 (3), 1164–1200.
- Manski, Charles F.**, "Identification of Endogenous Social Effects: The Reflection Problem," *The Review of Economic Studies*, 1993, 60 (3), 531–542.
- MESR**, "L'état de l'Enseignement Supérieur et de La Recherche," Technical Report, Ministère de l'Enseignement Supérieur et de la Recherche, Paris 2019.
- Mulhern, Christine**, "Beyond Teachers: Estimating Individual School Counselors' Effects on Educational Attainment," *American Economic Review*, 2023.

A. Institutional Background: Additional Details

A very clear overview of the French higher education landscape and its costs can be found in [Fack and Grenet \(2015\)](#). We summarise some of the important features of France's higher education below. Figure [A.1](#) provides an (somewhat simplified) illustration.

A.1. Access to Higher Education

In France, the only requirement to enter higher education is to obtain the end of high school exam, the *Baccalauréat* (hereafter Bac). Over 2013-2016, roughly 88% of students who took the Bac obtained it. Three types of Bac can be prepared by high school students, all of them corresponding to the type of high school track they are enrolled in. They are categories, in their more aggregate versions, as *general* (academic;), *technological* (technical), and *professional* (vocational). In 2021, half of Bac holders obtained a general Bac, the remaining half were divided between technological tracks (20%) and professional tracks. About three out of four high school students who obtained the Bac continued into tertiary education [MESR \(2019\)](#). This share is much higher for students from general and technical high school tracks as compared to students from vocational tracks.

A.2. Types of Higher Education Programs

The French higher education system is composed of five types of programs: non-selective public universities, (ii) selective *vocationally*-oriented post-secondary schools (Sections of Superior Technicians (*Sections de Techniciens Supérieurs (STS)*)), (iii) selective *technically*-oriented institutes (University Institutes of Technology (*Instituts Universitaires de Technologie (IUT)*)), (iv) selective *academically*-oriented preparatory classes (Preparatory Classes for the *Grandes Écoles (Classes Préparatoires aux Grandes Écoles (CPGE))*), and (v) other private schools (mostly engineering, business, art, and paramedical and social schools).

Selective institutions are free to select their applicants according to their own (undisclosed) criteria. Non-selective programs could not select its students. If capacity constraints were bindings, they distinguished applicants based on non-academic priority rules such as whether the student was from the same academic region as the institution and how applicants ranked the college-major in their rank-ordered list. Lotteries were implemented to break ties should capacity constraints continue to bind despite these priority criteria. See [Bechichi and Thebault \(2021\)](#) for more details, and analysis of these lotteries.

A.3. Cost of Higher Education

The cost of higher education depends exclusively on whether the institution is public or private. 82% of students are enrolled in a public institution. Public institutions charge annual tuition fees of slightly under 200 euros. There is no limit on the tuition fees private institutions can charge.

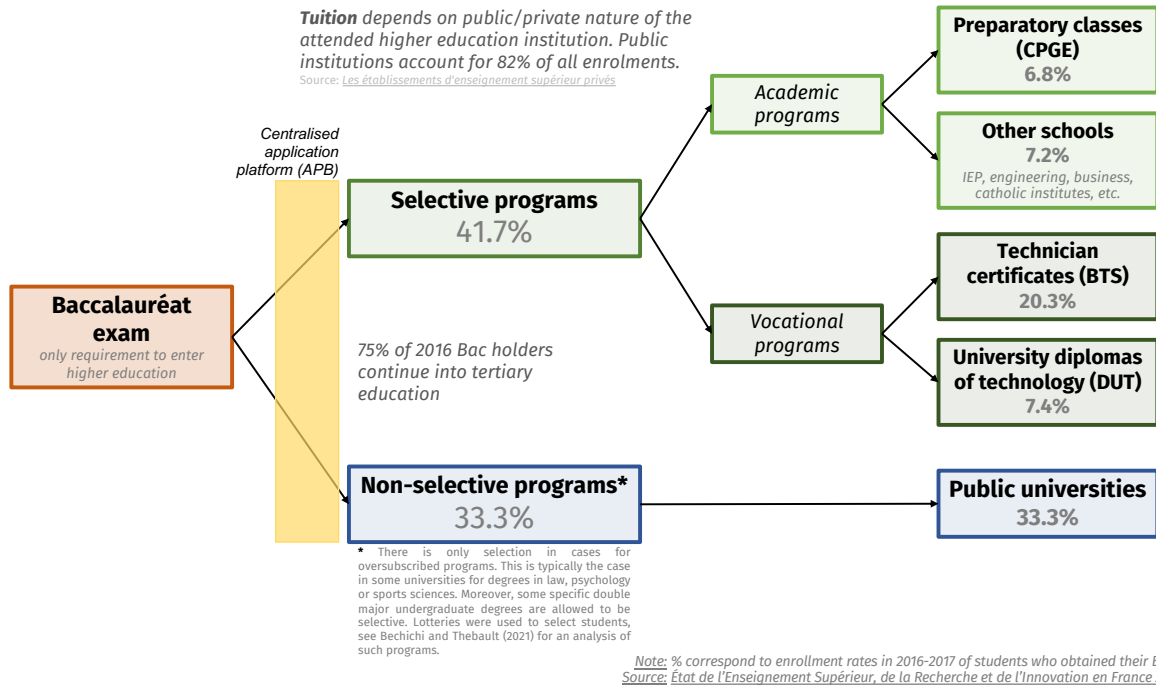


Figure A.1. Higher Education Landscape in France

A.4. Admission Post-Bac (APB)

From 2009 to 2017, students seeking admission to higher education programs were required to go through a centralised national platform called *Admission Post-Bac* (APB), where they could apply to both non-selective and selective programs. The APB system gathered roughly 12,000 programs and 800,000 applicants each year. Candidates submitting applications were asked to provide a rank-ordered list (ROL) of programs from January to March. Following this application phase, program administrators rank applicants. Each selective program produces its own specific ranking based on discretionary criteria and without any legal constraints. Selective programs are not required to rank all their applicants. The ranking for a non selective program is produced automatically by the centralised platform on the basis of applicants' non-academic priorities. In contrast to selective programs, a rank is assigned to all the applicants to a given non-selective program. It is important to note that the local decision rules or algorithms used by selective programs to rank applicants are not public information, neither for applicants nor for the centralised platform. The only information that the platform collects is the rank of each applicant, the outcome produced by local algorithms.

Taking into account programs' capacities in addition to applicants' ROL and programs rankings, applicants get offered a seat to their best feasible option through a three-rounds college-proposing deferred acceptance (DA) algorithm (Gale and Shapley, 1962) taking place from June to July (Figure 3.1). For each program j and each admission round k , applicant i gets offered a seat only if (i) the rank $r_{i,j}^k$ is above the cutoff c_j^k which corresponds to the rank of the last applicant receiving an offer; (ii) there is no higher-ranked

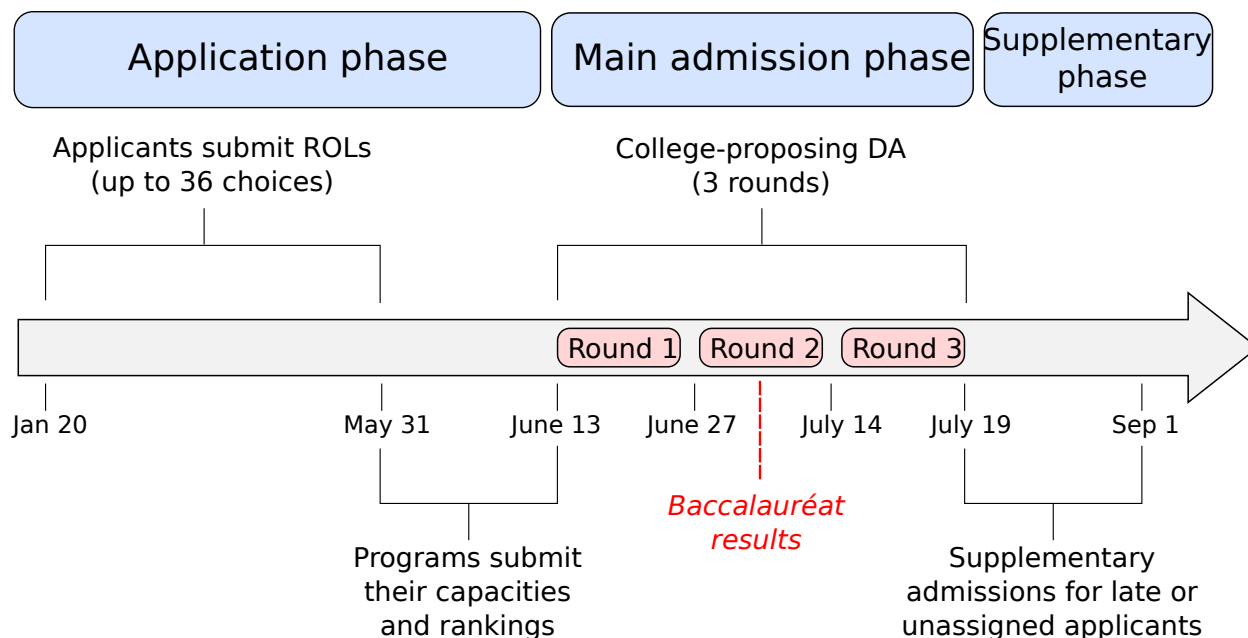


Figure A.2. Timeline of the Application and Admission Procedure into Higher Education Programs in France


program j' where applicant i is ranked above the cutoff $c_{j'}^k$. Applicants could accept the offer, turn it down or conditionally accept placement while waiting for applicants selected by higher-ranked programs to withdraw from the selection process in subsequent admission rounds. This sequential procedure implies that programs cutoffs could evolve from round 1 to round 3, always observing the following rule: $c_j^1 \leq c_j^2 \leq c_j^3$. The final results of the Bac were published between the second and the third rounds of the procedure. Students who failed the exam were not able to compete for a seat anymore, and their seats were re-offered in the third round. Finally, applicants could participate in supplementary rounds, which took place between June and September, and helped students to apply to programs with remaining seats. Figure 3.1 summarises the timeline of this process.

B. Running Variable: Additional Details

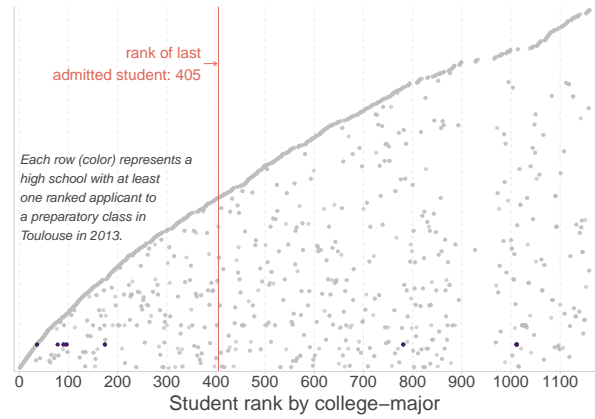
B.1. *Details on Running Variable*

In the main text we made a very slight simplification. In practice, college-majors rank their applicants within “*ranking groups*”. These ranking groups typically relate to applicants high school track, though they can be even more specific. This implies that the same college-major might have different rankings for different types of applicants. In our setting what matters for high school students is whether an older schoolmate was marginally admitted to a college-major or not, regardless of the ranking group to which they belonged. Thus, we overcome this minor issue (only 0.05% of high school x tracks, the level at which we conduct our analysis, have students in several ranking groups) by defining the high school’s best ranked applicant as the best ranked applicant to the college-major across all ranking groups. In the extremely few cases where the best rank is tied, we keep one at random.


B.2. Running Variable: An Illustration for a College-Major



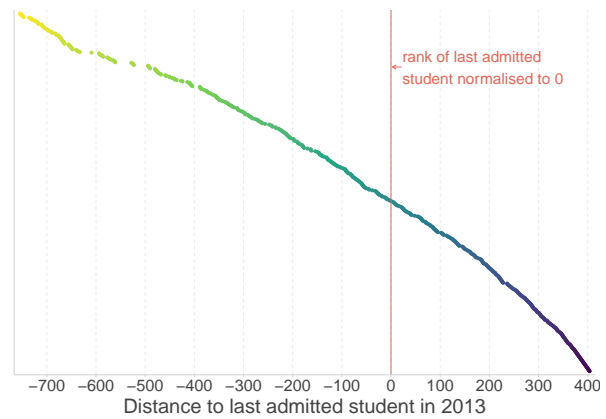
(a) All Applicants




(b) Applicants from one High School



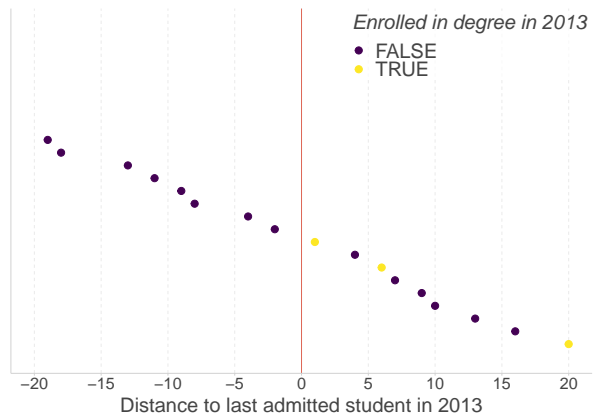
(c) Centering Around Rank of Last Admitted Student



(d) High School Defined by its Best Ranked Applicant



(e) Regression Discontinuity Bandwidth



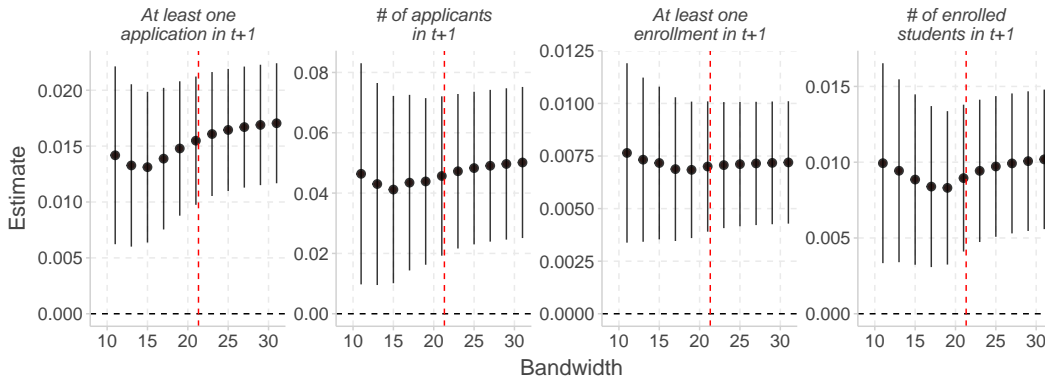
(f) Treatment Status

Figure B.1. An Example of the Running Variable for One College-Major


Notes: This figure shows how our running variable is computed for one college-major in 2013.

C. Robustness Checks


We assess the robustness of our baseline results to (i) varying the bandwidth over which the estimates are computed (Appendix Figure C.1, and (ii) estimating within-high school spillovers at placebo admission rank cutoffs (Appendix Figure C.2).



(a) College-Major Spillovers



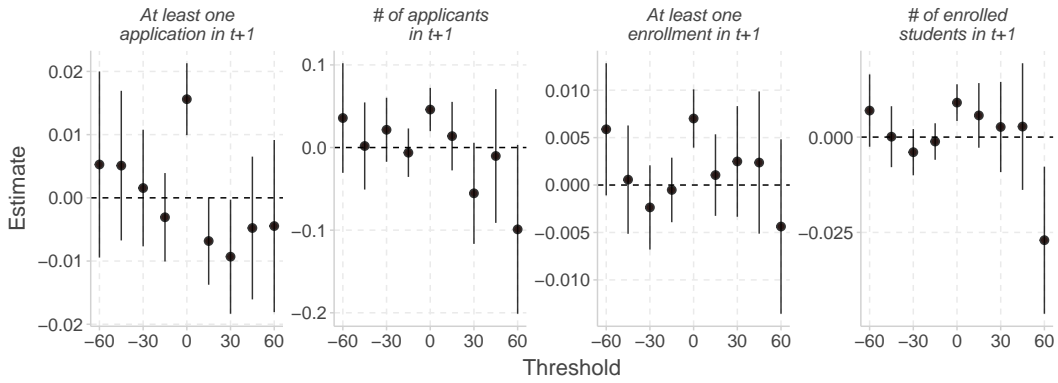
(b) College Spillovers



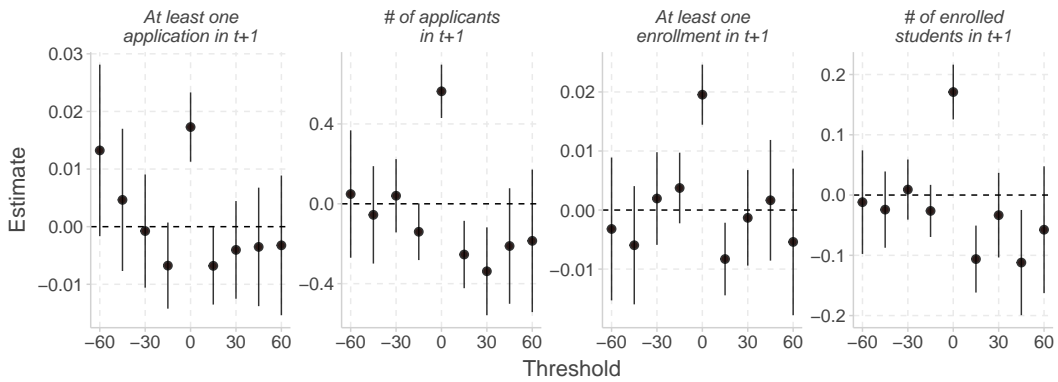
(c) Major Spillovers

Figure C.1. Robustness of Baseline Within-High School Spillovers to Using Different Bandwidths


Notes: This figure shows estimates of within-high school spillovers, varying the bandwidth over which the estimates are obtained. The application and enrollment outcomes are reported in the figure facet titles, while the type of spillover (college-major, college, or major) is indicated by the subfigure caption. The baseline bandwidth is denoted by the vertical dashed line. All the specifications in the figure correspond to local linear regressions using a triangular kernel, and include college-major - year fixed effects. Statistical significance is based on standard errors clustered at the high school - year level. Confidence intervals correspond to 95% confidence intervals.



(a) College-Major Spillovers



(b) College Spillovers



(c) Major Spillovers

Figure C.2. Robustness of Baseline Within-High School Spillovers to Placebo Admission Cutoffs

Notes: This figure shows estimates of within-high school spillovers, varying the admission cutoff at which the estimates are obtained. The application and enrollment outcomes are reported in the figure facet titles, while the type of spillover (college-major, college, or major) is indicated by the subfigure caption. All the specifications in the figure correspond to local linear regressions using a triangular kernel, and include college-major - year fixed effects. The bandwidth (21.32) corresponds to the smallest MSE-optimal bandwidth (Calonico et al., 2014) for the main college-major outcomes. Statistical significance is based on standard errors clustered at the high school - year level. Confidence intervals correspond to 95% confidence intervals.

D. Appendix Figures




Figure D.1. Exact Number of Enrolled Students in Treatment Year

Notes: This figure shows the number of enrolled students for a given high school and college-major as a function of the high school's distance to the last admitted student.

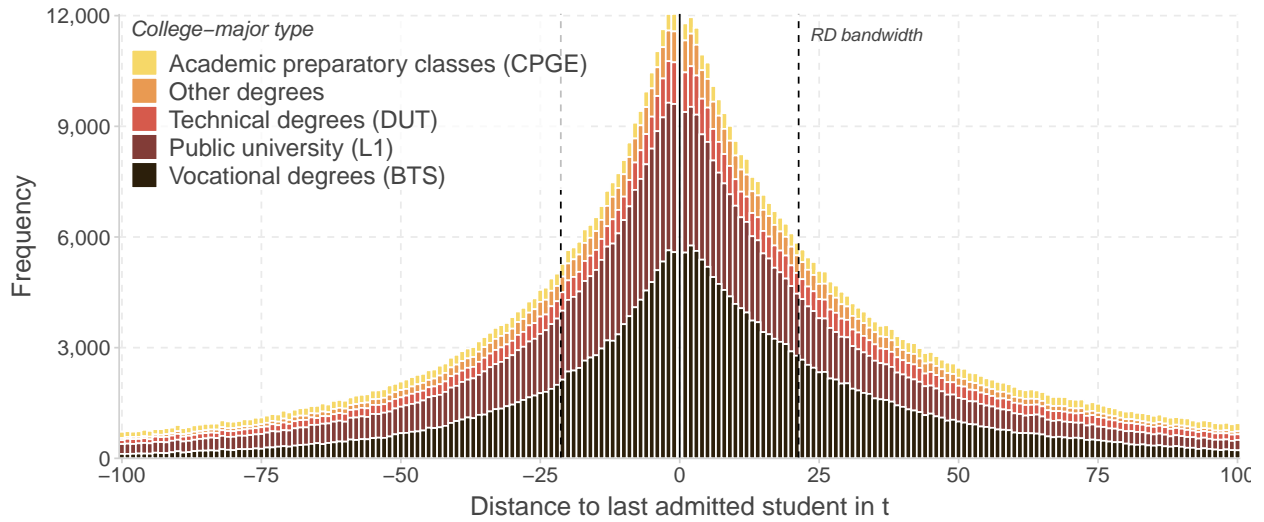


Figure D.2. Running Variable With Degree Type Composition

Notes: This figure shows the composition in terms of degree type of the running variable which corresponds to the rank of the high school's best ranked applicant by the college-major centered around the rank of the college-major's last admitted student. The dashed lines represent the the regression discontinuity (RD) bandwidth used in the analysis.

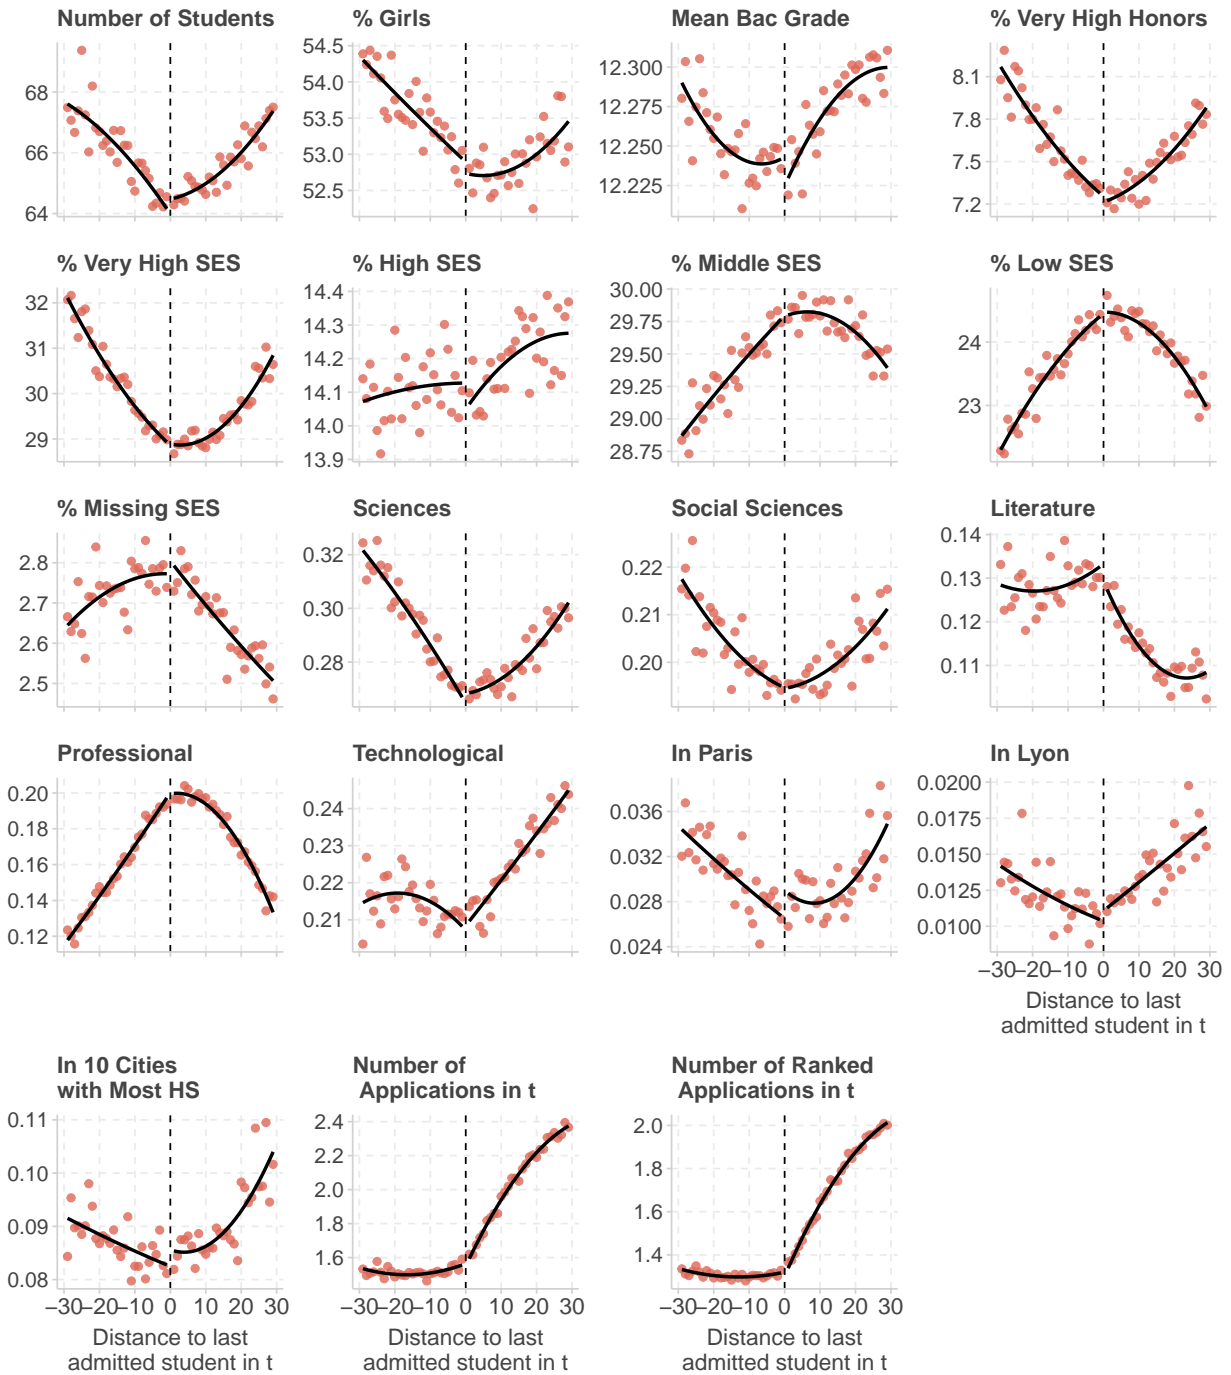


Figure D.3. Discontinuity in High School Characteristics

Notes: This figure shows non-parametric binned scatter plots of the relationship between various high school characteristics and distance to the last admitted student. The high school characteristics are reported in the figure facet titles. Each point corresponds to the average high school characteristic for high schools with distance to the last admitted student equal to the value on the x-axis. The fitted lines correspond to second-order polynomial fits through the conditional expectation.

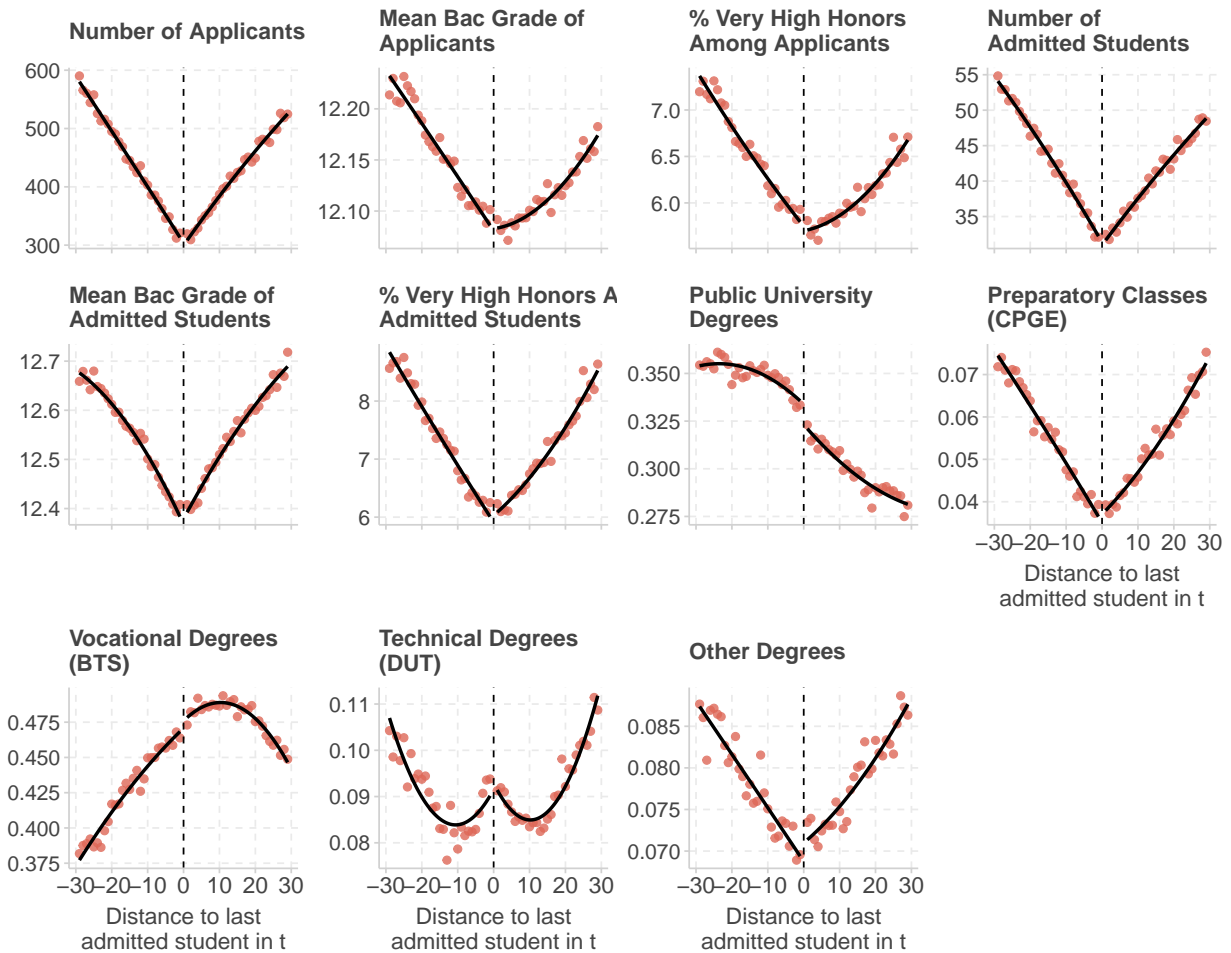
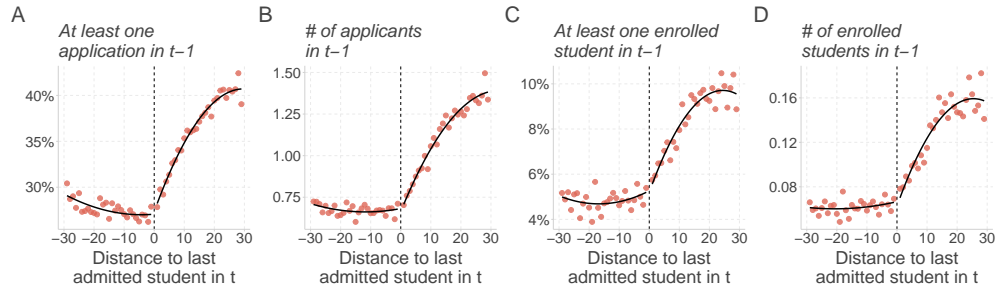
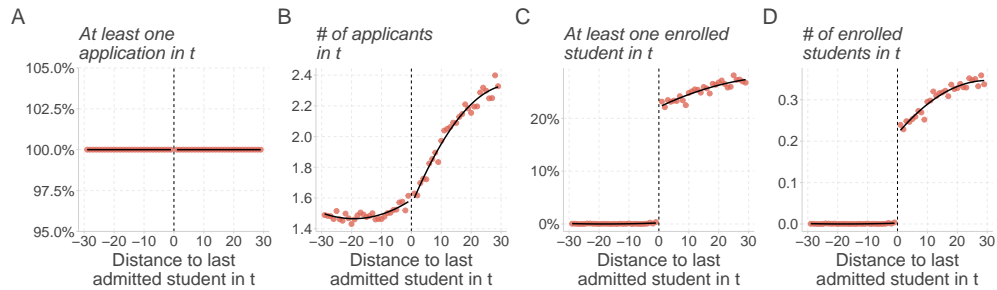


Figure D.4. Discontinuity in College-Major Characteristics

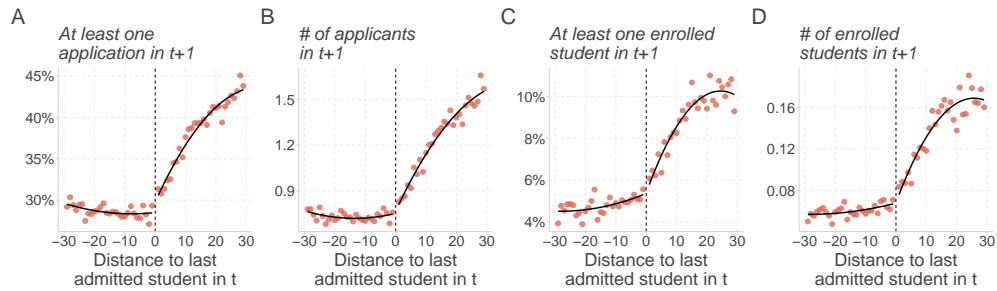
Notes: This figure shows non-parametric binned scatter plots of the relationship between various college-major characteristics and distance to the last admitted student. The college-major characteristics are reported in the figure facet titles. Each point corresponds to the average college-major characteristic for high schools with distance to the last admitted student equal to the value on the x-axis. The fitted lines correspond to second-order polynomial fits through the conditional expectation.



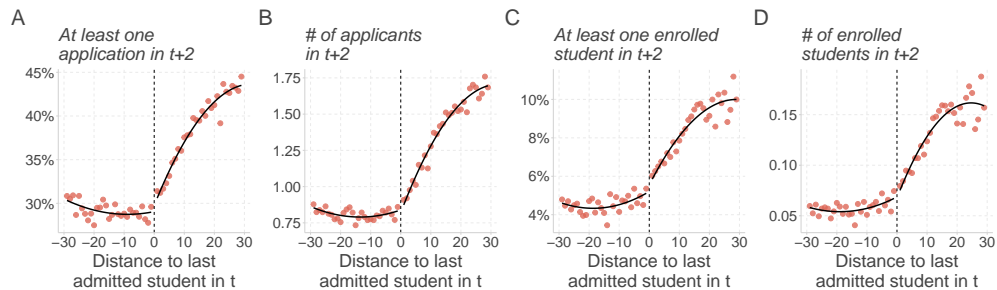
(a) Treatment Year - 1



(b) Treatment Year



(c) Treatment Year + 1



(d) Treatment Year + 2

Figure D.5. Event-Study Analysis Graphs

Notes: This figure shows non-parametric binned scatter plots of the relationship between high schools' application and enrollment outcomes for a college-major in different years and high schools' distance to the college-majors' last admitted student in t . "Treatment Year - 1" refers to the year prior to the older schoolmate's marginal admission, "Treatment Year" refers to the year of an older schoolmate's marginal admission, and "Treatment Year + 1" and "Treatment Year + 2" correspond, respectively to high schools' application and enrollment outcomes one and two years following the marginal admission of one of its students. The sample is restricted to treatment years 2014 and 2015 to ensure the sample is constant across estimates. Each point corresponds to the average college-major characteristic for high schools with distance to the last admitted student equal to the value on the x-axis. The fitted lines correspond to second-order polynomial fits through the conditional expectation.

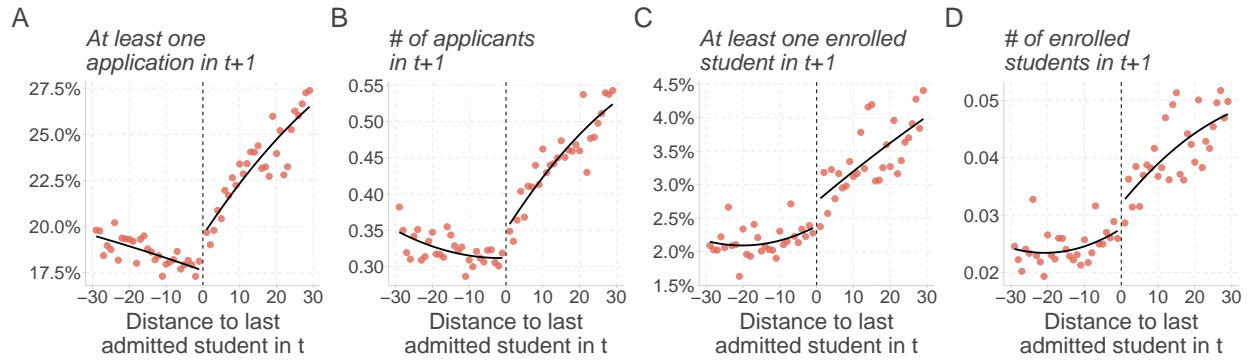


Figure D.6. Effects for College-Majors with No Applicants from High School in $t - 1$

Notes: This figure shows non-parametric binned scatter plots of the relationship between high schools' application and enrollment outcomes for a college-major in $t + 1$ and high schools' distance to the college-majors' last admitted student in t , for the subset of college-majors for which there were no applicants in a high school's $t - 1$ cohort of students. The application and enrollment outcomes are reported in the figure facet title. Each point corresponds to the average outcome value for high schools with distance to the last admitted student equal to the value on the x-axis. The fitted lines correspond to second-order polynomial fits through the conditional expectation.

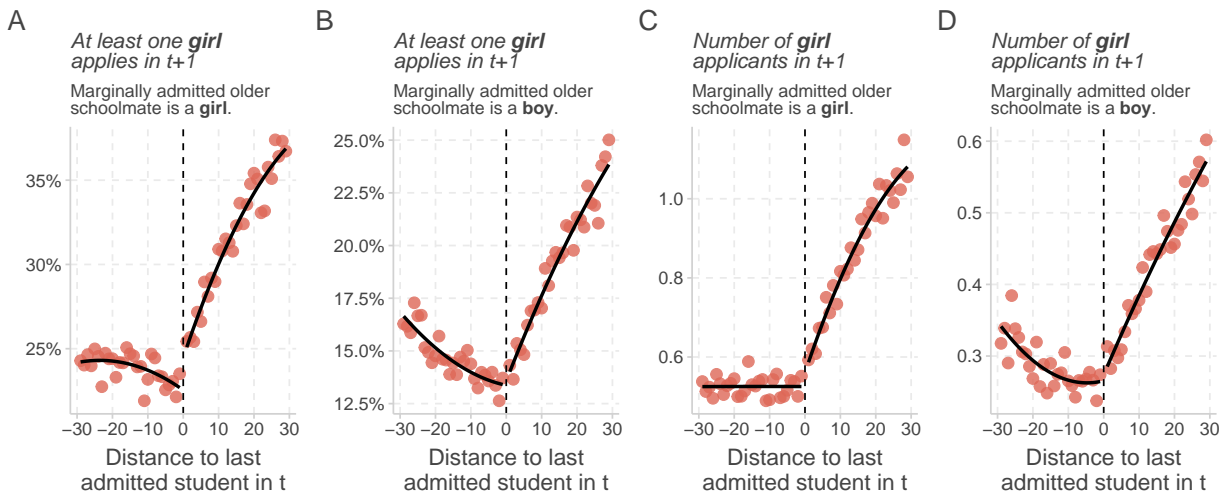


Figure D.7. Role Model Effects for Girls

Notes: This figure shows non-parametric binned scatter plots of the relationship between girls' application and enrollment outcomes for a college-major in $t + 1$ and high schools' distance to the college-majors' last admitted girl or boy in t . The application and enrollment outcomes are reported in the figure facet title. Each point corresponds to the average outcome value for girls in high schools with distance to the last admitted student equal to the value on the x-axis. The fitted lines correspond to second-order polynomial fits through the conditional expectation.

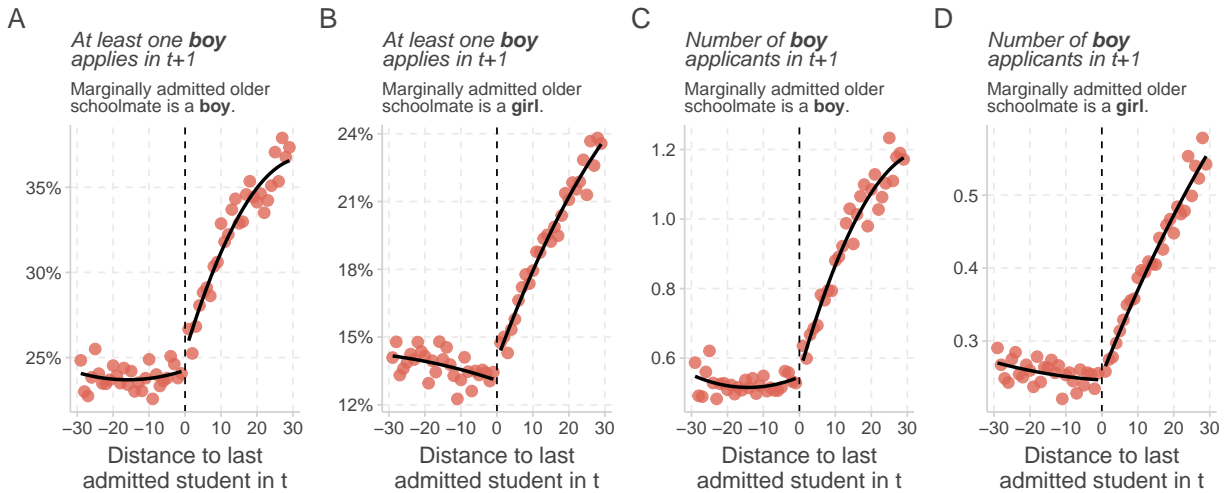


Figure D.8. Role Model Effects for Boys

Notes: This figure shows non-parametric binned scatter plots of the relationship between boys' application and enrollment outcomes for a college-major in $t + 1$ and high schools' distance to the college-majors' last admitted boy or girl in t . The application and enrollment outcomes are reported in the figure facet title. Each point corresponds to the average outcome value for boys in high schools with distance to the last admitted student equal to the value on the x-axis. The fitted lines correspond to second-order polynomial fits through the conditional expectation.

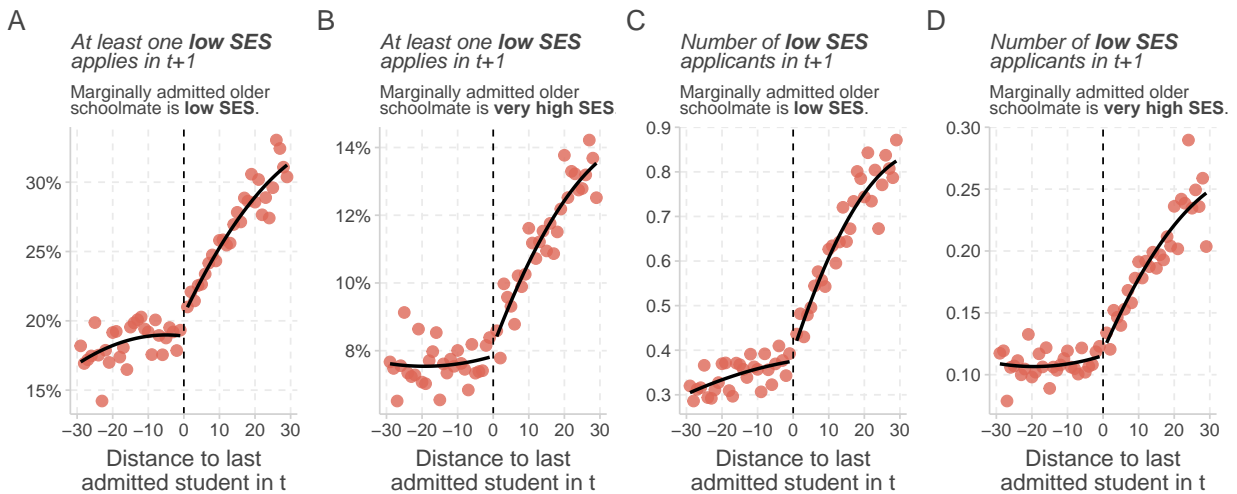


Figure D.9. Role Model Effects for Low SES Students

Notes: This figure shows non-parametric binned scatter plots of the relationship between low SES students' application and enrollment outcomes for a college-major in $t + 1$ and high schools' distance to the college-majors' last admitted low SES or very high SES student in t . The application and enrollment outcomes are reported in the figure facet title. Each point corresponds to the average outcome value for low SES students in high schools with distance to the last admitted student equal to the value on the x-axis. The fitted lines correspond to second-order polynomial fits through the conditional expectation.




Figure D.10. Role Model Effects for Very High SES Students

Notes: This figure shows non-parametric binned scatter plots of the relationship between very high SES students' application and enrollment outcomes for a college-major in $t + 1$ and high schools' distance to the college-majors' last admitted very high SES or low SES student in t . The application and enrollment outcomes are reported in the figure facet title. Each point corresponds to the average outcome value for very high SES students in high schools with distance to the last admitted student equal to the value on the x-axis. The fitted lines correspond to second-order polynomial fits through the conditional expectation.

E. Appendix Tables

Table E.1: Number of Observations at Each Sample Restriction

Restriction	Nb. College-Majors	% Change	Nb. High Schools	% Change	Nb. Total Obs.	% Change
Raw number	40,789	100%	53,159	100%	4,531,125	100%
+ At least one applicant ranked after last admitted student	34,975	85.75%	53,048	99.79%	4,167,954	91.98%
+ High schools with at least one applicant in two consecutive years	34,975	100%	51,376	96.85%	4,140,973	99.35%
+ No change in reported capacity between admission rounds	28,863	82.52%	51,150	99.56%	3,105,764	75%
+ Symmetrization of running variable	26,705	92.52%	49,714	97.19%	1,002,812	32.29%
+ Drop marginal student	26,209	98.14%	49,640	99.85%	987,688	98.49%
+ At least 2 obs. on both sides of cutoff within bandwidth	18,543	70.75%	49,434	99.59%	906,324	91.76%

Notes: This table shows the number of college-major - years, high school - years, and observations at each sample restriction mentioned in Section 3.2. High schools refer to high school x tracks.

3

High-Achieving, Low-Income Students and Higher Education: What Can Financial Aid Achieve?

Abstract

Why do high-achieving, low-income students enrol in higher education at lower rates and in lower quality institutions than their high-income peers? This paper assesses whether increased financial assistance can mitigate these gaps. Specifically, I estimate the impact of automatically granting additional financial support to high-achieving, low-income students who enrol in higher education. Using comprehensive administrative data for France and a regression discontinuity design, I find this policy had no significant effect on enrollment, persistence, graduation or academic performance in higher education, and small positive effects on geographic mobility. I also find no evidence that this aid induced eligible students to enrol in or switch to higher quality degrees during their studies. These null results do not appear to be driven by (i) students being unaware of this aid, (ii) crowding out of parents' financial support, or (iii) the aid being awarded on top of other grants. This highlights potential complementarities between financial aid and academic ability.

1. Introduction

GRADUATING from higher education provides one of the highest returns on investment an individual can make, especially when attending a selective institution (Bleemer, 2021; Black et al., 2023; Chetty et al., 2023). Yet, high-achieving, low-income students enrol at lower rates than their high-income peers, and when they do, they tend to attend lower quality institutions (Hoxby and Avery, 2013; Crawford et al., 2016; Dynarski et al., 2021; Hakimov et al., 2022; Campbell et al., 2022). This *undermatching* leads to large efficiency losses which could potentially be remediated by policy.

Understanding the factors underlying these gaps is therefore crucial to design effective policy responses. Are high-achieving, low-income students less aware of the benefits of attending higher education, and specifically selective institutions? Do they lack information about relevant programs or simply do not have the self-confidence to apply? Or is it that they require additional financial resources to attend these selective colleges? If the former reasons prevail, then informational/motivational interventions should be favored. However, if financial constraints are the dominant explanation, then targeted financial support would be the preferred policy.

In this paper, I analyse whether additional financial aid can serve as an effective way of inducing high-achieving, low-income students to pursue higher education and enrol in high-quality institutions, as well as persist and graduate in a timely manner. Specifically, I estimate the effects of a national financial aid scheme, the *aide au mérite*, introduced in 2008 in France, which automatically granted an additional 1,800 euros annually, for 3 years at most (the duration of a bachelor's degree), to eligible students who enrolled in a higher education institution. The only criteria to be eligible to the *aide au mérite* were that the student (i) be eligible to the national need-based grant program, and (ii) score at least 16 out of 20 (i.e. in the top 4.7% of exam takers) at the French end of high school exam, the *Baccalauréat* (henceforth Bac).

The targeted population of students thus corresponds very closely to [Hoxby and Avery \(2013\)](#)'s definition of high-achieving, low-income students (top 4% of U.S. high school students, and in bottom parental income quartile). By design, the *aide au mérite* was awarded on top of need-based grants which included a tuition fee waiver and annual cash allowances up to 5,500 euros for the most disadvantaged students. As such the *aide au mérite* represented at least a 40% top up in monthly allowances, a sizable increase in financial support.

Using administrative data on the universe of students obtaining the Bac between 2009 and 2014, I exploit the sharp discontinuity in eligibility to the *aide au mérite* at the 16/20 Bac grade threshold in a regression discontinuity design. This enables me to estimate the causal effect of eligibility to this additional financial aid in the Bac year on enrollment, degree quality, persistence, graduation and academic performance in higher education as well as geographic mobility.

I find that being eligible to the *aide au mérite* in the Bac year had precisely estimated zero effects on enrollment, persistence or graduation from higher education. For most outcomes, I can reject effects as small as one to three percentage points. In this context,

the enrollment margin is not particularly informative since, conditional on being eligible to a need-based grant, the enrollment rate around the 16 threshold is 94%. Moreover, students only become aware of their eligibility to the aide au mérite in July when Bac grades are released, which may limit the potential impact on enrollment. However, as in the U.S., persistence in higher education is a major concern in France. Around the 16 threshold, less than three out of four need-based grant eligible students are enrolled on time in 2nd year, and only just over half enrol in 3rd year on time. Thus, the null effects on persistence and graduation cannot be explained by students' late awareness of eligibility.

Additionally, I find no evidence that eligibility to the additional financial aid had an effect on the type or quality (proxied by the median Bac grade of students contemporaneously enrolling in the degree) of degree pursued. The null effects on degree quality remain for the degree enrolled in one year and two years later. This result rules out the hypothesis that eligible students become aware of the merit aid too late in the initial enrollment process but once aware subsequently choose to change tracks towards more selective degrees located in more expensive cities.

There is no discernible impact on other measures of higher education involvement such as the number of years enrolled in higher education or the highest level of study attained, nor on proxies for academic performance such as the likelihood of enrolling in a selective masters degree or the quality of the masters degree (again proxied by the median Bac grade of contemporaneous peers in the degree). Though I cannot observe students' undergraduate grades directly, this is indicative that academic performance does not appear to have been much influenced by eligibility to the aide au mérite. There is no clear sign of heterogeneous effects by gender or socio-economic background, suggesting these findings reflect true null effects and not heterogeneous effects that average out. This implies that high-achieving students' trajectories in higher education, even when they come from disadvantaged backgrounds, seem to be largely unaffected by the amount of financial support they receive. I do find evidence of positive effects on geographic location (Paris, and largest French cities) though the magnitude of the estimates are sensitive to the chosen bandwidth.

I exploit heterogeneity across specific subgroups to investigate three potential mechanisms that may underlie these null effects: (i) lack of information about eligibility to the aid, (ii) crowding out of parents' financial assistance, and (iii) the aid being awarded on top of need-based grants.

First, I find no evidence that the non-effect on *enrollment* might be driven by students

being unaware of the policy. Since the aid was automatically granted to eligible students (conditional on enrolling in higher education), take up is not a concern. However, the aide au mérite was introduced at the same time as a vast reform of the need-based grants system and therefore may not have been as salient to students as this latter change. Yet, the estimates are not larger for more recent Bac cohorts who are very likely to have been more aware of the policy, nor are they larger for students with more eligible high-school peers. This suggests that information deficits about the aide au mérite are unlikely to explain the null effect found on enrollment, though as discussed previously it could potentially be explained by students only becoming aware of eligibility late in the process. Since eligible students receive the financial aid once enrolled, there is no informational concerns for outcomes other than initial enrollment.

Second, I estimate the effects for students from the lowest-income families, who receive the highest need-based grants amounts but whose families are able to give them less than the amount of the aide au mérite on average (Grobon and Wolff, 2022). Thus, for these students, even if parental assistance is fully crowded out by the aide au mérite, they would still on net be better off financially. Admittedly, this will not necessarily be the case for students whose parents' give them more than 200 euros monthly, and for whom the aide au mérite could theoretically be fully compensated by crowding out. I find no effects for the lowest-income students, suggesting that the overall null effects are unlikely to be the result of crowding out of parental financial contributions fully compensating the amount received from the aide au mérite. I cannot rule out potential interactions between eligibility and parent income that may not go through the crowding out channel, though one could expect that if there were any effects for a subgroup of students they would most likely be for the most disadvantaged students.

Lastly, I observe no evidence that students who are eligible only to the tuition fee waiver and no cash allowance as part of their need-based grant exhibit greater behavioral responses to eligibility to the aide au mérite than students who are eligible to more generous monthly cash allowances as part of their need-based grants. These results hold even when restricting to students with very similar parent incomes, suggesting these differences are not simply the result of differential parent incomes. This implies that the null effects are likely not completely driven by the aide au mérite being awarded on top of other financial aid, thus limiting its potential ability to have any effect.

This mechanism analysis indicates that the most likely explanation for the lack of observed effects is that high-achieving, low-income students are not marginal students, in the sense that their higher education outcomes are not contingent on the amount of finan-

cial aid they are eligible to. This is in line with a number of studies who consistently find that the impact of financial aid on higher education outcomes tends to be small (or null) for the highest ability students while effects for lower ability students are sizable (Goodman, 2008; Cohodes and Goodman, 2014; Fack and Grenet, 2015; Bettinger et al., 2019; Angrist et al., 2022). These findings highlight potential complementarities between financial aid and academic ability. A fruitful future research avenue would be to investigate more precisely how the effects of financial aid vary along the student ability distribution.

I address two potential concerns with the empirical analysis: (i) the possibility for manipulation of students' Bac grade, and (ii) whether the purely symbolic "Very Good" honors associated with obtaining 16/20 at the Bac may contaminate the results.

First, the identification strategy relies on Bac grade not being manipulable. Since it corresponds to the weighted average of subject-level exam grades, manipulation by students is unfeasible. However, there is bunching of students just above the aide au mérite eligibility grade threshold due to review juries discretionarily increasing the grade of students close to this cutoff. To overcome this issue, I implement a *donut* regression discontinuity design, a commonly used solution in such instances (Barreca et al., 2016; Canaan and Mouganie, 2018; Angrist et al., 2019; Barr et al., 2022). Specifically, I drop students whose Bac grade is in the plausible range of discretionary adjustment, such that the observable characteristics of students are well-balanced around the cutoff.

Second, the symbolic honors associated with obtaining at least 16/20 at the Bac (*Mention Très Bien*) could have a direct impact on students' outcomes, for example through a psychological boost of getting this honor or because a very small number of higher education institutions may have special admission tracks for such high-achieving students. I find there is no discontinuity in outcomes at the 16 threshold for students *not* eligible to the need-based grant, rejecting this possible threat to identification.

This paper contributes to two distinct literatures. First, it speaks to the literature on high-achieving, low-income students initiated by Hoxby and Avery (2013). They documented that these students, despite being in top 5% of high school students' academic achievement, applied to significantly lower quality high education institutions than their high-income peers. This *undermatching* phenomenon has also been found in England (Campbell et al., 2022) and France (Hakimov et al., 2022). Many papers have tried to better understand its roots. For example, targeted and timely information about and mentoring on college options, application process and financial aid seem to mitigate part of the

gap (Hoxby and Turner, 2013; Carrell and Sacerdote, 2017). Certainty over the amount of financial aid received also significantly increases high-achieving, low-income students' enrollment rates (Dynarski et al., 2021; Burland et al., 2023). I contribute to this literature by showing that additional financial aid, without any individualised contextual information, is unlikely to reduce undermatching.

Second, my paper contributes to the vast literature on postsecondary financial aid¹, and specifically programs combining need- and merit-based components. I am able to study the effects of a national financial aid scheme on students in the top 5% of the academic ability distribution in a context with little uncertainty about the cost of higher education. The closest studies are Cohodes and Goodman (2014) (Massachusetts Adams Scholarship) and Andrews et al. (2020) (UT-Austin Longhorn Opportunity Scholars) who estimate the effects of financial aid with merit-based criteria, though in both cases the targeted population of students is more academically diverse than the aide au mérite (top 25% and top 30% respectively). Moreover, this existing evidence is based on U.S. state-level programs that are only available to students attending the state's flagship public universities. As such, these programs' effects are largely dependent on the quality of the state's higher education institutions, as shown by Cohodes and Goodman (2014), limiting their external validity. Conversely, the aide au mérite is a national-level program, covering all higher education institutions, and is therefore not contaminated by the effects of college quality by design.

Additionally, this paper sheds light on the potential mechanisms underlying the effects of financial aid, and points towards the importance of the academic level of the targeted student population. Though several studies have found that financial aid effects tend to be larger for low-ability students compared to high-ability students (Goodman, 2008; Cohodes and Goodman, 2014; Fack and Grenet, 2015; Bettinger et al., 2019; Angrist et al., 2022), this study puts much greater emphasis on this aspect as a key ingredient to the potential effectiveness of financial aid schemes. Lastly, the postsecondary financial aid literature is overwhelmingly U.S.-centred and analyses short-, medium- and long-run outcomes of financial aid schemes that, by and large, are meant to cover expensive tuition fees (exceptions are Fack and Grenet (2015) (France, need-based grant scheme), Dearden et al. (2014) and Murphy and Wyness (2022) (England/UK), Montalbán (forthcoming) (Spain), Baumgartner and Steiner (2006) (Germany) and Vergolini et al. (2014) (Province of Trento, Italy)). Much less is known about the impact of financial aid in higher education systems where tuition fees are significantly lower, where there may be centralised

¹See Herbaut and Geven (2020) for a systematic review of the evidence.

higher education application systems, and where financial aid is designed to cover living costs rather than tuition fees. This paper aims to help fill this gap in the literature.

The rest of the paper is organised as follows. Section 2 provides institutional background. Section 3 describes the data and sample used for the analysis, while Section 4 details the empirical strategy I adopt. Section 5 presents the main results and robustness checks, and Section 6 investigates the potential mechanisms. Section 7 concludes.

2. Institutional Background

2.1. *Aide au Mérite*

Overview.² The aide au mérite was introduced as part of a broader reform of higher education financial aid in France which came into effect in the 2008 academic year. This new grant consisted in nine monthly installments of 200 euros³, for a yearly total of 1,800 euros, for at most three years (the duration of a typical bachelor's degree in France) over the course of the student's undergraduate studies.⁴ Eligible students had to fulfill various academic requirements in order to continue receiving the aid (such as not failing their exams unless it was due to serious medical reasons, attending classes and exams) though it is unclear how scrupulously these were enforced in practice.⁵

Eligibility criteria. There were two eligibility criteria to the aide au mérite: (i) being eligible to the national need-based grant program⁶, and (ii) obtaining 16 out of 20 or above at the Bac, the French end of high school exam.⁷ I explain the most relevant aspects of these

²All details regarding the aide au mérite can be found in the [circulaire N°2008-1013 du 12 juin 2008](#) (in French).

³The amount was halved starting in the fall of 2015, with the reduced amount only applying to new recipients. My analysis is limited to the cohorts that benefited from the pre-reduction amount.

⁴This three-year limitation applied to students with a linear trajectory as well as to students who changed degree over the course of their studies. The only exception to this three-year limitation was for students in medical degrees who could benefit from this aid during the entirety of their medical studies.

⁵There were two exceptions to these requirements: (i) for first-year medical students, and (ii) for second-year preparatory class (*classes préparatoires aux grandes écoles*) students, who could repeat the grade without losing eligibility. These exceptions reflect the specificities of these programs in France: (i) there is very strict selection into second-year of medical studies due to a *numerus clausus*, and (ii) after the second year of preparatory classes, students take competitive exams in order to get into "elite" *Grandes Écoles*, and can choose to retake them the following year.

⁶Students whose parents did not pay any income tax were also eligible though in practice such cases appear to be extremely rare.

⁷The official criteria is actually to have obtained the highest honors (*Mention Très Bien*) at the Bac, which

two criteria below. Eligibility was automatically assessed each year based on students' Bac grade and their need-based grant status. Note that since eligibility to need-based grants can vary from year to year, a student could potentially be eligible to the aide au mérite in a given year and not be in the next. Students *receive* the grant amount only if they actually enrol in a higher education institution.

Annual quotas. Each academic region was annually allocated a given number of aide au mérite grants they could award to eligible students in their geographic purview. In practice, these quotas were not very binding: around 5% of students are registered in the data as not receiving the aide au mérite even they though they fulfil the required criteria and are enrolled in higher education.⁸

Bac. The Bac (abbreviation for *Baccalauréat*) is the French end of high school exam. It is organised in June each year. It consists in a multitude of subject-level exams (between 5 and over 15 depending on the track). A final grade out of 20 is then computed as a weighted average of subject grades. I refer to this final grade as Bac grade. Students scoring 10 or above obtain the Bac.

2.2. *Need-Based Grants*

Overview. The main higher education financial aid program in France is a need-based grants system called the *bourses sur critères sociaux*. In 2009-10, around 565,000 students benefited from such grants, representing roughly a third of students enrolled in higher education (MESR, 2011). A detailed analysis of these grants can be found in Fack and Grenet (2015).

Eligibility criteria. Eligibility to these need-based grants is assessed every year (regardless of eligibility status in the previous year) based on the combination of two criteria: (i) financial resources (parents' total gross income in year $n - 2$), and (ii) disadvantage points (up to 17; based on number of siblings and distance to the higher education institution).⁹

corresponds to obtaining at least 16/20. For ease of understanding, I use the latter formulation.

⁸It is unclear what rule academic regions used to allocate the grants among eligible students. In any case, since I conduct an intent-to-treat analysis, and as students cannot know in advance whether they will actually receive the aide au mérite or not, this is not a big concern for the analysis.

⁹Specifically, disadvantage points are awarded based on (i) the number of additional dependent children in the family (2 points per additional dependent child, 4 points per additional dependent child in higher education), and (ii) the distance to the higher education institution from the student's home address (30-249

Importantly, students have to file an online application with the above information in order for their eligibility to be assessed by the higher education financial aid agency.

Amounts. Each combination of parent income and disadvantage points corresponds to a given *echelon* of financial aid, which gives right to an amount of cash allowance handed out in ten monthly installments (September to June). Appendix Table B.1 shows the combinations of (parent income, disadvantage points) and the related echelon for the academic year 2009-10. Between 2009-10 and 2012-13 there were 7 echelons, from 0 (least generous) to 6 (most generous). In 2013 two additional echelons were created, "0 bis" (between 0 and 1) and 7 (most generous). Appendix Table B.2 displays the annual amounts of aid given to each echelon between 2009 and 2014. Echelon 0 students were only exempted from paying tuition and student social security fees and did not receive any cash allowance while echelon 6 students received just over 4,000 euros (in addition to being exempt from tuition and student social security fees).

Amount of aide au mérite discussion. The amount of the aide au mérite, 200 euros per month over 9 months, may seem like a small amount in absolute terms, yet it was actually quite generous in relative terms. First, since aide au mérite recipients were exempt from paying tuition fees (which are relatively low in the first place, less than 200 euros per year), unlike most financial aid in the U.S. this grant aimed to help cover living expenses. Second, the aide au mérite was very generous relative to the need-based grants these students received: 125% and 43% of the minimum and maximum need-based grant amounts respectively for the 2009-10 academic year. Lastly, it represented about a third of the average student's monthly budget, estimated to be around 700 euros by (Fack and Grenet, 2015).

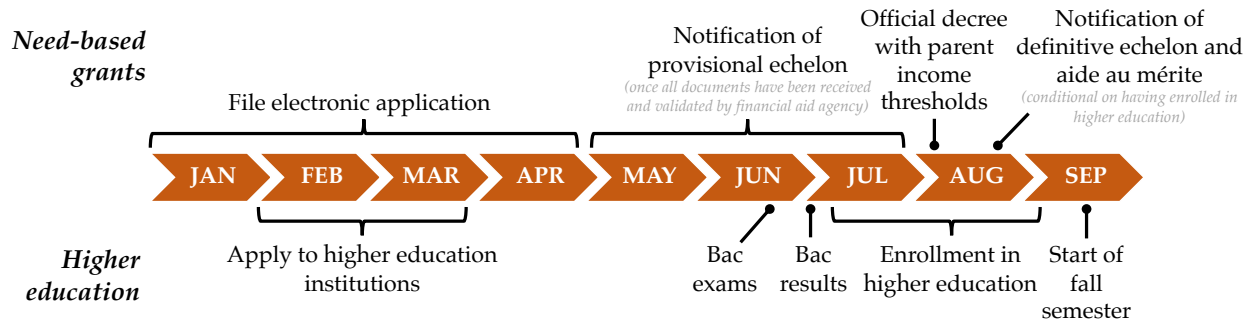
2.3. Timeline

A summary of the timeline of events is presented in Figure 3.1.

Need-based grants. Students apply electronically for need-based grants between January and April/May. They can apply after this deadline if circumstances justify it. The financial aid agency processes applications to ensure all supporting documents have been transmitted and are in due form. Students are then informed of their provisional need-

km = 1 point, ≥ 250 km = 2 points).

Figure 3.1. Timeline of Events



Notes: This figure shows the timing of events related to need-based grants, to higher education, and to the Bac, for the months between January and September in a typical year. Exact dates vary slightly year to year.

based grant echelon (should there be one).¹⁰ The official parent income thresholds for eligibility to each need-based grant echelon are published between mid-July and mid-August. Students receive a definitive notification of their need-based grant echelon and aide au mérite amounts once they officially enrol in higher education.

Higher education. Students register their applications to postsecondary degrees by the end of January through a centralised platform called *Admission Post-Bac* (APB) or directly to the higher education institutions not on the platform. They receive a decision on their applications in various waves between June and mid-July. They officially enrol over the summer months.

Bac. High school students take the exams of the Bac in June and get their Bac grade in early July.

2.4. Higher Education in France

Structure. A very clear overview of the French higher education landscape and its costs can be found in [Fack and Grenet \(2015\)](#). I only describe the key institutional elements needed to understand the analysis here.

High school students wishing to pursue postsecondary education essentially have the

¹⁰This provisional notice also includes eligibility to the aide au mérite in years other than the Bac year, since in the Bac year the Bac grade is unknown at this stage of the process.

choice between five pathways: (i) non-selective public universities¹¹, (ii) selective *vocationally-oriented* post-secondary schools (*Sections de Technicien Supérieur* (STS)), (iii) selective *technically-oriented* institutes (*Instituts Universitaires de Technologie* (IUT)), (iv) selective *academically-oriented* preparatory classes (*Classes Préparatoires aux Grandes Écoles* (CPGE), also known as *prépas*), and (v) other selective private schools.¹² The only criteria to continue into higher education in France is to obtain the Bac. The Bac grade obtained does not play any role in one's likelihood of being accepted in a selective degree except in extremely few instances.¹³ Among students who obtained the Bac in 2009, 78% enrolled in a higher education institution in the same year (MESR, 2011). 44% were enrolled in a public university, 25% in a STS, 11% in a IUT, 10% in a CPGE, and 10% in other private institutions.

Cost. The cost of higher education in France varies depending on the type of institution attended. Annual tuition fees at public universities and IUT are set at very low levels (171 euros in 2009-10 at the undergraduate level; in addition students paid a social security contribution of 198 euros). The cost of studying in a STS or CPGE depends on whether the institution is public (no tuition fees, only student social security fee) or private. Fees at private schools are very heterogeneous and can go up to several thousand euros per year. The main financial barrier concerns living costs rather than tuition fees. Fack and Grenet (2015) estimate, using data from 2010, that the total average budget for a nine-months academic year is around 6,300 euros, i.e. 700 euros per month. As such, available financial aid is insufficient to fully cover these expenses, requiring parents to help out if they can and students to work on the side of their studies. The French Ministry of Employment estimates that on average between 2013 and 2015 23% of students enrolled in higher education were employed at some point during their studies, of which 33% in a job not linked to their studies and not only over the summer (DARES, 2017).

¹¹The vast majority of public universities' undergraduate degrees were not selective, other than having obtained the Bac. There is selection only in instances where there are more applicants to the degree than available seats, though this selection was done through a random lottery. In practice, this concerns few degrees. See Bechichi and Thebault (2021) for additional details.

¹²Mostly engineering and business schools as well as institutions not attached to a university (accounting, architecture, ...), art schools, and paramedical and social schools.

¹³A known exception is Sciences Po Paris which for a long time admitted students with an exceptionally high Bac grade.

3. Data, Sample Construction and Descriptive Statistics

3.1. Data

I combine data from four administrative sources provided by the statistical offices of the French Ministry of Education (MENJS-DEPP) and the Ministry of Higher Education (MESRI-SIES) using a unique anonymised student identifier:

OCEAN (*Organisation des Concours et des Examens Académiques et Nationaux*), 2006-2020. Covers the universe of high school students taking the Bac. For each student, the dataset provides information on the high school (location, type), the Bac track, the Bac grade, as well as socio-demographic characteristics such as age, gender, and socio-economic status (SES) defined by the Ministry of Education based on the legal guardian's occupation (see [Bonneau et al. \(2021, p.72\)](#) for the detailed classification).

AGLAE (*Application pour la Gestion du Logement et de l'Aide à l'Étudiant*), 2008-2018. Covers applications for higher education public grants. It contains information on which type of aid students applied for, whether they obtained the grant, if rejected what the reason was, if accepted what the echelon was, parent income, number of disadvantage points, and whether the student *received* the aide au mérite¹⁴.

SISE (*Système d'Information sur le Suivi de l'Étudiant*), 2008-2020. Covers almost all students enrolled and graduating from a higher education institution other than vocational tracks (STS) and academic preparatory classes (CPGE). For each student, it contains information on the higher education institution and degree enrolled in, the year in the degree, the length of the degree, and whether the degree has been obtained or not.

BPBAC (*Base Post-Bac*), 2009-2020. Covers students enrolled in a vocational track or in an academic preparatory class and contains the same information as SISE except it does not contain information on graduation.

Coverage. As some paramedical and social diplomas as well as some artistic and cultural

¹⁴Note that the data itself does not indicate whether the student was *eligible* to the aide au mérite. I infer eligibility status based on Bac grade and need-based grant eligibility.

higher education institutions are not covered by SISE, [Bonneau et al. \(2021\)](#) estimate that for the 2016-17 academic year around 90% of students in higher education were covered by the SISE and BPBAC data.

3.2. *Sample*

I restrict my sample to high school students who (i) obtained the Bac between 2009 and 2014, (ii) had a unique and non-missing student identifier, (iii) obtained the Bac only once over the 2006-2014 period¹⁵, (iv) did not have a missing Bac grade¹⁶, and (v) were eligible to a need-based grant in their Bac year. The reason for restriction (i) is that the BPBAC data for the 2008-09 academic year is missing students' identifiers, and the amount of the aide au mérite was halved for students entering higher education in 2015. The final sample contains 1,101,658 students.¹⁷

3.3. *Descriptive Statistics*

Appendix Table [B.4](#) provides some descriptive statistics for three samples: (i) the full sample, (ii) aide au mérite eligibles in their Bac year, and (iii) the sample of students scoring between 15 and 17 at the Bac, which will be used in the empirical analysis. Out of the roughly 1 million students in the sample, about 55,000 were eligible to the aide au mérite in their Bac year. Only 5% of the sample obtained above 16/20 at the Bac, a necessary condition to be eligible to the aide au mérite. This proportion matches very closely [Hoxby and Avery \(2013\)](#)'s percentage of "high-achieving" students (top 4%¹⁸ of all U.S. high school students). Compared to the full sample, aide au mérite eligibles are slightly more likely to be female, and to come from higher income and SES families. They are also more concentrated among the lower echelons of need-based grants reflecting their less disadvantaged socio-economic backgrounds. They are also significantly more likely to have favorable higher education outcomes, including obtaining a degree. Reflecting the French higher education system, aide au mérite eligibles are much more likely to enrol in an academic preparatory class or other private schools relative to the entire sample.

¹⁵I make this restriction to drop students who may have strategically obtained the Bac again in order to obtain above 16 and receive the aide au mérite. In practice, extremely few students obtain the Bac more than once over this period (0.34%).

¹⁶Only 0.1% of students satisfying (i)-(iii) have a missing Bac grade.

¹⁷See Appendix Table [B.3](#) for the sample size at each additional restriction.

¹⁸Specifically, students "who score at or above the 90th percentile on the ACT comprehensive or the SAT I (math and verbal) and who have a high school grade point average of A- or above."

4. Empirical Strategy

I use a regression discontinuity design to estimate the causal effect of being eligible to the aide au mérite in the Bac year on various higher education outcomes such as enrollment, degree quality, persistence and graduation.¹⁹ Specifically, I exploit the eligibility discontinuity at 16/20 at the Bac: need-based grant eligible students scoring at or above this threshold are automatically eligible to the aide au mérite, while students scoring just below are not. Estimating an OLS regression of the outcome on a dummy variable for being eligible to the aide au mérite would yield a biased estimate because eligible students have higher grades than non-eligible students, which is correlated with better higher education outcomes. On either side of the threshold, students should be very similar and differ only with respect to their eligibility to the aide au mérite (Imbens and Lemieux, 2008; Lee and Lemieux, 2010; Cattaneo et al., 2019).

Importantly, this analysis therefore estimates *intent-to-treat* effects, i.e., the effect of being *eligible* to the aide au mérite *in the Bac year*, since only students who eventually enrol in higher education actually *receive* this aid. Because of the endogeneity of any subsequent outcome, even when estimating the impact on an outcome measured in the years following the Bac year, I continue to compare students who were at the margin of being eligible to the aide au mérite *in their Bac year*, versus those who were at the margin of not being eligible *in their Bac year*.²⁰

4.1. Estimation Details

Running variable. The running variable is the student's Bac grade, which I denote $Bac\ grade_i$, where i refers to a student. This grade lies between 8 and 20. If $Bac\ grade_i$ is greater than or equal to 16, then the student is eligible to the aide au mérite, otherwise the student is not.

Local average treatment effect. Any discontinuity in higher education outcomes between students around the 16 threshold can be interpreted as the causal effect of being eligible

¹⁹One may want to also conduct a difference-in-differences analysis by comparing need-based grant eligible students below and above 16/20 at the Bac before and after the introduction of the aide au mérite in 2008. However, simultaneously to the introduction of the aide au mérite, a vast reform of need-based grants was implemented, simplifying the disadvantage points calculation from 8 criteria to only 2 (see [circulaire n°2007-066 du 20 mars 2007](#) (in French) for details on the pre-2008 system). Moreover, the BPBAC data for 2008 is missing student identifiers further complicating such an analysis.

²⁰Appendix Figures A.1 and A.2 show the evolution of students in the full sample's aide au mérite and need-based grant status over time. These are helpful to better interpret the ITT estimates.

to the aide au mérite in the Bac year. Given student i 's outcome y_i , the causal effect is identified by:

$$\beta^{RDD} = \lim_{\varepsilon \rightarrow 16^+} \mathbb{E}(y_i | \text{Bac grade}_i = \varepsilon) - \lim_{\varepsilon \rightarrow 16^-} \mathbb{E}(y_i | \text{Bac grade}_i = \varepsilon) \quad (3.1)$$

β^{RDD} is the causal effect of being eligible to the aide au mérite in the Bac year on the outcome of interest for students who obtained a Bac grade very close to 16.

Main specification. My main specification for estimating the causal effect of the aide au mérite is as follows:

$$y_i = \alpha + \beta(\text{Bac grade}_i - 16) + \gamma \text{Aide au mérite}_i + \lambda(\text{Bac grade}_i - 16) \times \text{Aide au mérite}_i + \theta X_i + \varepsilon_i, \quad (3.2)$$

where y_i is student i 's higher education outcome regressed on i 's *Bac grade*, *Aide au mérite* is an indicator for aide au mérite eligibility ($\text{Bac grade}_i \geq 16$), the interaction between both variables, and in some specification a rich vector of pre-treatment control variables X_i (gender, age, SES, high-school track and Bac cohort). The coefficient of interest is γ . Adding the control variables is not needed for identification but they can improve the estimates' efficiency (Calonico et al., 2019). ε_i is the error term.

Following Cattaneo et al. (2019)'s guidelines, the coefficient of interest is estimated non-parametrically using local linear regressions. Specifically, linear regressions are fit on both sides of the threshold using a triangular kernel which gives more weight to observations near the threshold. I report all estimates using two bandwidths: (i) the mean squared error (MSE) optimal bandwidths computed using Calonico et al. (2014)'s procedure and which differ across specifications, and (ii) the (15, 17) bandwidth which has the advantage of keeping the sample constant across outcomes.

Inference. Inference for the MSE-optimal bandwidths is based on Calonico et al. (2014)'s' robust bias-corrected procedure, which corrects for estimated bias in the point estimate to construct the confidence interval. As such the reported robust 95% confidence intervals are not necessarily centered around the point estimate (but around the point estimate plus the estimated bias; see Cattaneo et al. (2019) for more details). For the (15, 17) bandwidth estimates, I report conventional confidence intervals.

Identifying assumption. The main identifying assumption underpinning the regression discontinuity design is that at the limit of the 16/20 threshold students scoring just below are essentially identical to those scoring just above. This assumption would be threatened if students attempted to score exactly at 16 or just above.

Since the Bac grade is computed as a weighted average of individual subject grades, there is little scope for grade “manipulation” (i.e. aiming for an exact grade) by students. Therefore, though the aide au mérite may possibly have incentivised students to obtain higher grades at the Bac, it is highly implausible that it may have led them to obtain a grade just above 16. However, as evident from Figure 3.1, which displays the distribution of Bac grades for the full sample, there is bunching around important Bac grade cutoffs²¹, and in particular around 16/20.

The reason for these very sharp discontinuities in the Bac grade distribution is that once all subject exams have been graded, jurys review students’ grades and can discretionarily slightly increase the grades of students close to these important thresholds in order for these students to obtain a final grade just above the relevant threshold. The decision to “upgrade” a student is not based on any rule and is entirely left to the discretion of the members of the jury in charge of the student’s file, which is composed of their grades at the Bac and comments by their professors. This upgrading of original grades by jurys poses an important threat to the identification strategy since adjusted students differ from non-adjusted ones (more on this below), along margins that are likely related to outcomes.


To overcome the non-random upgrading of students’ grades, I adopt a *donut* regression discontinuity strategy, which consists in dropping observations near the cutoff which have potentially been manipulated (Barreca et al., 2016). This is a very common method used in cases where there might be non-random heaping in the running variable (e.g., Angrist et al. (2019), Barr et al. (2022)).²² In particular, it is used by Canaan and Mouganie (2018) to exploit the 10/20 Bac obtention grade threshold to estimate the returns to higher-education quality for low ability students.

It is impossible to precisely identify which students have been upgraded and which have not. From Appendix Figure A.4, it is clear that students are not adjusted above 16.05.

²¹Obtaining at least 10 implies the student obtains the Bac, at 12, 14 and 16 students are awarded various honours called *mention*, respectively mention *Assez Bien* (Quite Good), mention *Bien* (Good) and mention *Très Bien* (Very Good). Note that this bunching is not specific to need-based grant eligible students as can be seen in Appendix Figure A.3.

²²Other notable examples include, birth weight: Bharadwaj et al. (2013), high school GPA: Cohodes and Goodman (2014), blood alcohol content: Hansen (2015), Maimonides rule: Angrist et al. (2019), age-based disability program: Deshpande et al. (2021), among others.

Figure 3.1. Distribution of Bac Grades, 2009-2014



Notes: This figure shows the distribution of Bac grades of full sample students, in panel (a) between 8 and 20, and in panel (b) between 15 and 17. For panel (a) each bin represents the number of students who obtained a Bac grade in $[X, X + 0.1)$, while in panel (b) each bin represents the number of students who obtained a Bac grade in $[X, X + 0.05)$. The full sample consists in students from the 2009-2014 Bac cohorts who obtained the Bac, filed a financial aid application in their Bac year and were eligible to a need-based grant in their Bac year.

Thus the upper limit for the donut can be reasonably set to 16.05 (included). To select the lower limit of the donut, I estimate discontinuities in observable characteristics (see next subsection) for lower limits from 15.6 to 15.95 in .05 increments (the upper limit is fixed at 16.05). The results are presented in Appendix Figure A.5. The smallest donut boundaries which balances characteristics around the 16 threshold is $[15.7, 16.05]$. Thus, in my donut specification I drop observations between 15.7 (included) and 16.05 (included) from the regressions.

Moreover, the 16/20 Bac grade is associated with a symbolic highest honors (*Mention Très Bien*), which could have a direct impact on students' outcomes. This effect could be driven by the psychological boost of getting this honor or because some higher education institutions may have special admission tracks for such students.²³ I show in the following subsection that there is no discontinuity in outcomes at the 16 threshold for students *not* eligible to the need-based grant in the Bac year (i.e., who are not eligible to the aide au mérite), suggesting the associated honors at the threshold does not threaten the validity of the identification strategy.

Robustness. In Appendix C, I assess the robustness of the main results to (i) estimating

²³A policy called "*Dispositif Meilleurs Bacheliers*" (Best Graduates Rule), which guaranteed a seat in a selective degree to students scoring in the top 10% of their high school at the Bac, was introduced in 2014. This affects only the last Bac cohort of my sample. In any case, very few students actually benefited from the program (900 in 2017), and is not based specifically on the 16/20 threshold.

equation (3.2) using a second-order polynomial of the running variable, (ii) including numerous pre-treatment student characteristics as controls, and (iii) varying the size of the bandwidth used for point estimation (Appendix Figures A.10-A.11). Overall none of these robustness checks affect the main conclusions.

4.2. *Tests of Design Validity*

I conduct three validity checks of my empirical design. Specifically I test for discontinuities (i) in students' pre-treatment observable characteristics, (ii) in predicted outcomes based on pre-treatment observables, and (iii) at various placebo grade thresholds. For completeness, I conduct all three validity tests (i) without any observation exclusions ("No Donut") and with the donut specification ("Donut [15.7, 16.05]"), and (ii) using both the MSE-optimal and (15, 17) bandwidths.

Discontinuity in observable characteristics. Table 3.1 reports estimates of equation (3.2) where the left-hand side variable is indicated in the first column's rows. Appendix Figure A.6 presents these estimates graphically. Each row therefore corresponds to a different regression. While many characteristics are statistically and economically significant in the full sample specifications, almost none remain significant and the coefficients are small in magnitude in the donut specification. Coherently, the only characteristics that are large in magnitude and statistically significant in the no donut specifications are characteristics that are observed by review jurors, such as gender, age, academic region, and Bac track. Importantly, there is no discontinuity in the donut specification in terms of SES, parent income²⁴, and echelon level, all characteristics likely to correlate with higher education outcomes.

Discontinuity in outcome prediction. My second validity test consists in estimating a simple prediction model of the outcomes under study using as predictors students' characteristics in Table 3.1 and then estimating equation (3.2) with the prediction on the left-hand side. The prediction model includes no interactions and is estimated by OLS. Appendix Table B.5 reports the results of this exercise. Consistent with the previous test on observables, all the estimates are statistically significant in the no donut specification while no estimate is significant in the donut case, though some predictions have admittedly low adjusted R^2 .

²⁴For parent income, I drop 202 unreasonably large or negative observations ($\geq 100,000$ and $\leq -100,000$).

Table 3.1: Discontinuity Estimates for Pre-Treatment Observable Characteristics

	Mean <i>[15.5, 15.7]</i>	No Donut		Donut [15.7, 16.05]		
		MSE-Optimal	<i>Bandwidth:</i>		MSE-Optimal	(15, 17)
			(15, 17)	(15, 17)		
	(1)	(2)	(3)	(4)	(5)	
<i>Demographic</i>						
Female	0.58	0.181*** [0.13, 0.25]	0.072*** [0.05, 0.09]	0.026 [-0.03, 0.1]	0.026** [0, 0.05]	
Age	18.09	-0.106*** [-0.16, -0.07]	-0.035*** [-0.06, -0.01]	0.02 [-0.02, 0.07]	0.014 [-0.02, 0.05]	
French Nationality	0.98	-0.008* [-0.02, 0]	-0.007*** [-0.01, 0]	-0.003 [-0.02, 0.01]	-0.004 [-0.01, 0]	
<i>Parent SES</i>						
Very High SES	0.24	-0.002 [-0.03, 0.02]	0.001 [-0.01, 0.02]	0.019* [0, 0.04]	0.018 [0, 0.04]	
High SES	0.17	0.004 [-0.01, 0.02]	0.004 [-0.01, 0.02]	-0.014 [-0.04, 0.01]	-0.013 [-0.03, 0.01]	
Middle SES	0.3	-0.006 [-0.03, 0.02]	-0.002 [-0.02, 0.01]	0.002 [-0.02, 0.03]	0.005 [-0.02, 0.03]	
Low SES	0.26	-0.003 [-0.03, 0.03]	0.001 [-0.02, 0.02]	0.036* [0, 0.08]	-0.002 [-0.02, 0.02]	
Missing SES	0.03	0 [-0.01, 0.01]	-0.004 [-0.01, 0]	-0.03*** [-0.06, -0.01]	-0.007 [-0.02, 0]	
Parent Income	26,928	-627* [-1502, 52]	-496* [-999, 7]	-196 [-943, 467]	-403 [-1102, 295]	
<i>Need-Based Grants</i>						
Echelon 0-0bis	0.34	-0.01 [-0.03, 0.01]	-0.011 [-0.03, 0.01]	-0.001 [-0.03, 0.02]	-0.005 [-0.03, 0.02]	
Echelon 1	0.19	0.001 [-0.02, 0.02]	0.001 [-0.01, 0.02]	-0.003 [-0.03, 0.03]	0.004 [-0.02, 0.02]	
Echelon 2-4	0.24	0.008 [-0.01, 0.03]	0.008 [-0.01, 0.02]	-0.003 [-0.03, 0.03]	-0.003 [-0.02, 0.02]	
Echelon 5-7	0.23	0.001 [-0.02, 0.02]	0.001 [-0.01, 0.02]	0.002 [-0.02, 0.03]	0.004 [-0.02, 0.03]	
<i>Geographic</i>						
Paris Academie	0.02	-0.001 [-0.01, 0.01]	-0.002 [-0.01, 0]	0.003 [-0.01, 0.02]	0.002 [-0.01, 0.01]	
5 Largest Academies	0.29	0.057*** [0.03, 0.1]	0.034*** [0.02, 0.05]	0.025 [-0.02, 0.08]	0.012 [-0.01, 0.04]	
<i>High-School</i>						
General Track	0.79	0.199*** [0.16, 0.25]	0.033*** [0.02, 0.05]	0.003 [-0.04, 0.04]	-0.002 [-0.02, 0.02]	
Technological Track	0.12	0.019** [0, 0.04]	0.017*** [0.01, 0.03]	0.02 [-0.02, 0.06]	0.005 [-0.01, 0.02]	
Professional Track	0.09	-0.037*** [-0.05, -0.03]	-0.05*** [-0.06, -0.04]	-0.006 [-0.02, 0.02]	-0.003 [-0.02, 0.01]	
Private	0.2	-0.026* [-0.06, 0]	-0.021*** [-0.04, -0.01]	-0.033 [-0.09, 0.01]	-0.013 [-0.03, 0.01]	

Notes: This table reports estimates of the discontinuity in student characteristics at the aide au mérite eligibility threshold (16/20 Bac grade). The student characteristics are reported in the first column's rows. For example, column (2) indicates that, using the MSE-optimal bandwidth and keeping all observations, the estimated discontinuity in the share of female students around 16/20 is 18.1 percentage points, with the share of females among students with Bac grade in [15.5, 15.7] being 58%. This discontinuity estimate is 2.6 percentage points in the donut specification (column (4)), which excludes students with Bac grades in [15.7, 16.05]. The MSE-optimal bandwidth, obtained using the *rdrobust* R package, varies across each outcome. For MSE-optimal bandwidth estimates, the ranges in brackets correspond to associated robust 95% confidence intervals, while they correspond to associated conventional 95% confidence intervals for (15, 17) bandwidth estimates. Statistical significance is computed based on the relevant p-value and ***, **, and * indicate significance at 1, 5, and 10%, respectively.

Discontinuity at placebo thresholds. Lastly, I use three placebos to validate the donut approach and to ensure the results are not driven by the potential psychological effect or preferential admission from being awarded the highest honors at 16/20. Specifically, I estimate equation (3.2) at (i) grades 14/20, and (ii) 15/20, where there should be no effect since nothing specific happens at these grades, and (iii) at grade 16/20 for students *not* eligible to a need-based grant, who are therefore not eligible to the aide au mérite and for whom no effects should be found as well. Since there is no bunching at grade 15, the no donut estimates are also informative and should be small in magnitude.

Table 3.2 shows the estimates obtained for these placebo tests for the six main academic outcomes using the MSE-optimal bandwidth (the estimates using the (15, 17) bandwidth are reported in Appendix Table B.6). First, all coefficients for the grade 15 placebo are very small in magnitude and overwhelmingly insignificant in the no donut specification. This is reassuring since these coefficients should indeed be zero. Second, the estimates for other placebos in the no donut specification are sizable and statistically significant, which is expected considering grade adjustments are made based on students characteristics correlated with higher education outcomes. Third, all coefficients in the donut specifications are very small in magnitude and insignificant. This alleviates any concern that the results conflate the effect of the aide au mérite with other factors occurring at the 16/20 threshold as well.

Overall, these three validity tests suggest employing the donut specification strongly limits the potential bias induced by Bac grade adjustments.

5. Main Results

In this section I present the main results of the analysis. The educational outcomes analysed relate to (i) enrollment, (ii) degree quality, (iii) persistence in higher education, (iv) degree completion, and (v) academic performance. I also assess whether there are any effects of eligibility to the aide au mérite on geographic mobility. Table 3.1 summarises the results. For each outcome I report estimates using the MSE-optimal and the (15, 17) bandwidths, excluding observations between 15.7 and 16.05. All the detailed regression tables are relegated to Appendix C.

Table 3.2: Placebo Analysis

	Enrollment in Bac Year	Enrollment in 2nd Year in Bac Year + 1	Enrollment in 3rd Year in Bac Year + 2	Number of Years in Higher Education	Highest Level of Study Attained	Obtaining a Degree
	(1)	(2)	(3)	(4)	(5)	(6)
<i>Panel A. Grade 15</i>						
No Donut	0.009*** [0, 0.02]	0.005 [-0.01, 0.02]	0.006 [-0.01, 0.02]	0.034 [-0.03, 0.12]	0.034* [-0.01, 0.09]	0 [-0.01, 0.01]
Donut [14.7, 15.05]	-0.004 [-0.02, 0]	0 [-0.02, 0.03]	0.009 [-0.01, 0.03]	0.026 [-0.13, 0.13]	0.039 [-0.07, 0.11]	-0.001 [-0.03, 0.03]
<i>Panel B. Grade 14</i>						
No Donut	0.009*** [0, 0.02]	0.043*** [0.03, 0.06]	0.111*** [0.1, 0.13]	0.595*** [0.53, 0.69]	0.599*** [0.54, 0.68]	0.087*** [0.08, 0.1]
Donut [13.7, 14.05]	0.002 [-0.01, 0.01]	0 [-0.01, 0.02]	0.003 [-0.02, 0.02]	0.073 [-0.04, 0.16]	0.063* [-0.01, 0.14]	0.003 [-0.02, 0.02]
<i>Panel C. Grade 16 for Students Not Eligible to a Need-Based Grant in Bac Year</i>						
No Donut	0.083*** [0.07, 0.1]	0.094*** [0.08, 0.12]	0.103*** [0.08, 0.13]	1.19*** [1.04, 1.41]	0.991*** [0.87, 1.16]	0.176*** [0.14, 0.22]
Donut [15.7, 16.05]	-0.016 [-0.04, 0]	-0.008 [-0.03, 0.02]	-0.001 [-0.02, 0.02]	-0.028 [-0.25, 0.16]	-0.054 [-0.21, 0.07]	-0.005 [-0.04, 0.04]

Notes: This table reports estimates of the discontinuity in higher education outcomes for three placebo grade thresholds: grade 15 (panel A), grade 14 (Panel B), and grade 16 for students not eligible to a need-based grant (panel C). Each higher education outcome is reported in the column headers. The grade 15 and grade 14 placebos are estimated over the full sample. The grade 16 placebo is estimates on the sample of students from the same Bac cohorts who were not eligible to a need-based grant in their Bac year. In all three cases, students on both sides of placebo grade threshold are not eligible to different amounts of financial aid. I report full sample estimates (“No Donut”) and estimates obtained when excluding students with Bac grades in [placebo – 0.3, placebo + 0.05] (“Donut [...]”). I use the MSE-optimal bandwidth, with results using the (15, 17) bandwidth available in Appendix Table B.6. This bandwidth, obtained using the *rdrobust* R package, varies across each specification. Associated robust 95% confidence intervals are reported in brackets. For example, in the no donut specification, that is including all students from the full sample, the estimated discontinuity in enrollment in Bac year at grade 15 (first row) is 0.009 percentage points, with associated robust 95% confidence interval [0, 0.02]. When excluding students with Bac grade in [15.7, 16.05], the estimated discontinuity is -0.004 percentage points, with associated robust 95% confidence interval [-0.02, 0]. Statistical significance is computed based on the robust p-value and ***, **, and * indicate significance at 1, 5, and 10%, respectively.

5.1. Academic Outcomes: Enrollment, Degree Quality, Persistence, Graduation, and Academic Performance

Enrollment. I start by investigating the causal effect of eligibility to the aide au mérite in the Bac year on the probability of enrolling in higher education. Figures 3.1a and 3.1b display, respectively, the probability of being enrolled in higher education in the Bac year, and the probability of being enrolled in any year between 2009 and 2020, as a function of Bac grade. Each dot corresponds to a 0.05 grade bin²⁵. The black line is the local linear regression line with a triangular kernel applied to the (15, 17) bandwidth, excluding observations between 15.7 and 16.05 (grey dots).

Enrollment for students around the 16/20 cutoff is very high: conditional on being eligible to a need-based grant, the probability of enrolling in higher education in the Bac year is approximately 95%, while the probability of any enrollment is above 97%. There is a slightly increasing and linear relationship between these probabilities and Bac grade

²⁵Each bin includes the lower bound and excludes the upper bound.

Table 3.1: Effects of Eligibility to the aide au merite in Bac Year on Higher Education Outcomes

	Mean <i>[15.5, 15.7]</i>	Point Estimate	95% Confidence Interval	Bandwidth	# obs. left	# obs. right
	(1)	(2)	(3)	(4)	(5)	(6)
<i>Panel A. Enrollment</i>						
Enrollment in Bac Year	0.94	-0.005	[-0.02, 0.01]	(15.04, 16.96)	33,599	24,826
Enrollment in Bac Year	0.94	-0.005	[-0.02, 0.01]	(15, 17)	35,234	25,357
Any enrollment	0.97	-0.001	[-0.01, 0.01]	(14.89, 17.11)	43,505	27,783
Any enrollment	0.97	-0.001	[-0.01, 0.01]	(15, 17)	35,234	25,357
<i>Panel B. Persistence</i>						
Enrollment in 2nd Year in Bac Year + 1	0.73	-0.001	[-0.04, 0.02]	(15.09, 16.91)	30,331	23,948
Enrollment in 2nd Year in Bac Year + 1	0.73	-0.006	[-0.03, 0.02]	(15, 17)	35,234	25,357
Enrollment in 3rd Year in Bac Year + 2	0.56	0.015	[-0.02, 0.05]	(15.13, 16.87)	27,203	23,175
Enrollment in 3rd Year in Bac Year + 2	0.56	0.006	[-0.02, 0.03]	(15, 17)	35,234	25,357
Number of Years in Higher Education	5.23	-0.007	[-0.21, 0.15]	(15.19, 16.81)	24,015	21,613
Number of Years in Higher Education	5.23	-0.028	[-0.14, 0.09]	(15, 17)	35,234	25,357
Highest Level of Study Attained	4.29	-0.038	[-0.18, 0.07]	(15.23, 16.77)	21,646	20,890
Highest Level of Study Attained	4.29	-0.058	[-0.13, 0.02]	(15, 17)	35,234	25,357
<i>Panel C. Degree Quality</i>						
Median Bac Grade of Degree in Bac Year	13.28	-0.012	[-0.15, 0.08]	(15.12, 16.88)	26,671	22,174
Median Bac Grade of Degree in Bac Year	13.28	0.001	[-0.09, 0.09]	(15, 17)	33,000	24,211
Median Bac Grade of Degree in Bac Year + 1	13.33	0.011	[-0.17, 0.14]	(15.28, 16.72)	17,691	18,574
Median Bac Grade of Degree in Bac Year + 1	13.33	0.021	[-0.07, 0.11]	(15, 17)	31,843	23,778
Median Bac Grade of Degree in Bac Year + 2	13.39	-0.018	[-0.21, 0.15]	(15.32, 16.68)	14,062	16,772
Median Bac Grade of Degree in Bac Year + 2	13.39	0.019	[-0.07, 0.1]	(15, 17)	28,287	22,518
<i>Panel D. Degree Completion</i>						
Obtain a Degree	0.62	-0.022	[-0.05, 0]	(14.81, 17.19)	47,773	29,075
Obtain a Degree	0.62	-0.026**	[-0.05, 0]	(15, 17)	35,234	25,357
<i>Panel E. Academic Performance</i>						
Enrollment in a Masters Degree	0.69	-0.019	[-0.06, 0.01]	(15.22, 16.78)	22,335	20,989
Enrollment in a Masters Degree	0.69	-0.02*	[-0.04, 0]	(15, 17)	35,234	25,357
Enrollment in a Selective Masters Degree	0.23	-0.011	[-0.05, 0.01]	(15.12, 16.88)	28,440	23,236
Enrollment in a Selective Masters Degree	0.23	-0.009	[-0.03, 0.01]	(15, 17)	35,234	25,357
Median Bac Grade of Masters Degree	13.72	0.05	[-0.11, 0.18]	(15.21, 16.79)	15,163	16,160
Median Bac Grade of Masters Degree	13.72	0.058	[-0.03, 0.15]	(15, 17)	23,059	19,286

Notes: This table reports estimates of the discontinuity in higher education outcomes at the aide au mérite eligibility threshold (16/20 Bac grade), excluding students with Bac grades in [15.7, 16.05]. The higher education outcomes are reported in the first column's rows. Column (1) reports the mean outcome for students with Bac grade in [15.5, 15.7]. Estimates (col. (2)) are reported for both the MSE-optimal (upper row) and the (15, 17) (lower row) bandwidths. The MSE-optimal bandwidth, obtained using the *rdrobust* R package, varies across each outcome. For MSE-optimal bandwidth estimates, the associated confidence intervals (col. (3)) correspond to robust 95% confidence intervals, while they correspond to conventional 95% confidence intervals for (15, 17) bandwidth estimates. Column (4) reports the bandwidth over which the local linear regressions are estimated (it is centered around 16), while columns (5) and (6) report, respectively, the number of observations used for estimation to the left and right of the 16/20 threshold. For example, column (2) indicates that, using the MSE-optimal bandwidth (first row), the estimated discontinuity in enrollment in Bac year around 16/20 is -0.005 percentage points, with associated robust 95% confidence interval [-0.02, 0.01] ([15.5, 15.7] baseline: 94%). This estimate is also -0.005 percentage points in the (15, 17) bandwidth specification (second row). The MSE-optimal bandwidth for this specification is (15.04, 16.96). Statistical significance is computed based on the relevant p-value and ***, **, and * indicate significance at 1, 5, and 10%, respectively. All the detailed regression tables can be found in Appendix C.

though it is relatively mild. We can see graphically that (non-adjusted) students just below 16 have, on average, worse outcomes than their lower grade counterparts, though the sample size is very small. This is expected since review jurys choose to upgrade students

partly based on students' professor's comments and more unruly students will therefore be less likely to be upgraded.

There is no visual evidence of any discontinuity in enrollment at the 16/20 aide au mérite eligibility grade threshold. Somewhat reassuringly, students just above 16, who are overwhelmingly students who have been upgraded, have very similar outcomes to students a bit further above in the grade distribution. This suggests that unless one believes the effects of eligibility to the aide au mérite to be extremely local, the donut specification should capture reasonably well the causal effects had there been no adjustments.

The estimates reported in *Panel A* of Table 3.1 confirm the lack of effects seen in the figures. The estimates are insignificant, very small (-0.005 and -0.001 for both bandwidths) and precisely estimated. As shown in Appendix Tables C.1 and C.2 adding controls or using a second order polynomial of the running variable does not alter the estimates. These results imply that eligibility to the aide au mérite in the Bac year did not have any effect on enrollment in higher education.

Despite not seeming to have impacted enrollment behavior, eligibility to the aide au mérite may have induced students into medical degrees or academic preparatory classes, since in both cases students could continue benefiting from the aid even if they failed a year (see footnote 5 for details). Table 3.2 estimates discontinuities in the likelihood of being enrolled in a given type of degree (Appendix Figure A.7 shows these estimates graphically).²⁶ All coefficients are of very small magnitude and statistically insignificant, suggesting eligibility to the aide au mérite did not lead students to substitute towards degrees or institutions for which the aide au mérite might have given them an advantage.

Degree Quality. Second, I estimate whether eligibility to the aide au mérite in the Bac year affected the quality of degrees in which students enrolled in the Bac year up to 2 years after the Bac. I proxy degree quality by the median Bac grade of all students contemporaneously²⁷ enrolled in the degree. Degrees are defined at the higher education institution x major level (e.g., BSc Mathematics at Paris I). By construction, degree quality

²⁶Since 13% of students have multiple enrollments, I follow [Bonneau et al. \(2021\)](#) and assign a main enrollment to each student using the following priority rule: (i) engineering school, (ii) business school, (iii) Institutes of Political Studies (IEP), (iv) preparatory classes (CPGE), (v) professional vocational diploma (STS), (vi) technical vocational diploma (IUT), (vii) other private schools, (viii) public universities.

²⁷I opt for not defining degree quality as the median Bac grade of students enrolled in the *previous* year because (i) students in new degrees cannot be allocated a degree quality, (ii) students enrolled in preparatory classes (CPGE) and vocational tracks (STS) in 2009 cannot be allocated a degree quality due to the missing student identifiers for these degrees in 2008, and (iii) because a number of public universities merged in 2014 preventing me from allocating students in these merged institutions a degree quality.

Figure 3.1. Higher Education Outcomes as a Function of Bac Grade Around the Aide au Mérite Eligibility Threshold (16/20)



Notes: This figure shows non-parametric binned scatter plots of the relationship between various higher education outcomes and Bac grade, around the aide au mérite eligibility threshold (16/20 Bac grade). The higher education outcomes are reported in the subfigure captions. The sample used corresponds to students from the 2009-2014 Bac cohorts who obtained the Bac, filed a financial aid application and were eligible to a need-based grant. Each red bubble corresponds to the average outcome value for students with Bac grade in $[X, X + 0.05)$, with the size of the bubble corresponding to the number of observations in that grade range. The grey bubbles represent the excluded donut observations, that is observations in $[15.7, 16.05]$. The black fitted lines correspond to local linear regressions with a triangular kernel on each side of the threshold, using the (15, 17) bandwidth, and excluding donut observations. Note that each subfigure has its own y-axis scale.

Table 3.2: Effects of Eligibility to the Aide au Mérite in Bac Year on Degree Choice

	Mean [15.5, 15.7]	Bandwidth:			
		MSE-Optimal		(15, 17)	
		(2)	(3)	(4)	(5)
Public University (excl. Medical Degrees)	0.28	-0.003 [-0.04, 0.03]	-0.014 [-0.05, 0.01]	-0.016 [-0.04, 0.01]	-0.018 [-0.04, 0]
Medical Degrees	0.13	0.013* [0, 0.03]	0.012 [-0.01, 0.04]	0.014 [0, 0.03]	0.013 [0, 0.03]
Vocational Diploma (STS)	0.12	0.001 [-0.02, 0.03]	-0.001 [-0.02, 0.01]	-0.001 [-0.02, 0.02]	-0.001 [-0.01, 0.01]
Technical Diploma (IUT)	0.1	0.005 [-0.01, 0.03]	0.004 [-0.01, 0.03]	0.001 [-0.01, 0.02]	0.001 [-0.01, 0.02]
Academic Preparatory Classes (CPGE)	0.25	-0.016 [-0.06, 0.02]	-0.006 [-0.04, 0.02]	-0.003 [-0.03, 0.02]	-0.001 [-0.02, 0.02]
Other (Business and Engineering Schools, IEP)	0.06	0.001 [-0.04, 0.03]	0.001 [-0.04, 0.04]	0.001 [-0.01, 0.01]	0.001 [-0.01, 0.01]
Controls			✓		✓

Notes: This table reports estimates of the discontinuity in main enrollment degree in Bac year at the aide au mérite eligibility threshold (16/20 Bac grade), excluding students with Bac grades in [15.7, 16.05]. The enrollment degrees are reported in the first column's rows. Estimates for two different bandwidths are reported: the MSE-optimal and (15, 17) bandwidths. The MSE-optimal bandwidth, obtained using the *rdrobust* R package, varies across each outcome. For MSE-optimal bandwidth estimates, the ranges in brackets correspond to associated robust 95% confidence intervals, while they correspond to associated conventional 95% confidence intervals for (15, 17) bandwidth estimates. Control variables are gender, age, SES, Bac track, and Bac cohort. Column (1) reports the mean of the row enrollment type for students with Bac grade in [15.5, 15.7]. For example, column (2) indicates that, using the MSE-optimal bandwidth without controls, the estimated discontinuity in the main enrollment in Bac year being a public university degree (excluding medical degrees) around 16/20 is -0.3 percentage points, with the share of such main enrollments among students with Bac grade in [15.5, 15.7] being 28%. This discontinuity estimate is -1.6 percentage points when using the (15, 17) bandwidth (column (4)). Statistical significance is computed based on the relevant p-value and ***, **, and * indicate significance at 1, 5, and 10%, respectively.

can only be computed for students who actually enrol in a higher education institution. Focusing on degree quality is informed by the striking evidence from the U.S., England and France showing that high-achieving, low-income students apply to and enrol in lower quality degrees than their high-income peers (Hoxby and Avery, 2013; Campbell et al., 2022; Hakimov et al., 2022).

Figures 3.1g-3.1i graphically displays the relationship between degree quality in various years and Bac grade. For all degree quality measures, this relationship is linear and increasing. As for other outcomes, there is no clear discontinuity at 16 for any degree quality. Table 3.1 Panel D reports the associated regression discontinuity estimates. All estimates are close to zero and insignificant, suggesting being eligible to the aide au mérite in the Bac year has no effect on the quality of the degree in which one enrolls. I also find no

effects on degree quality in subsequent years. This rules out the hypothesis that eligible student might have opted to switch to higher quality degrees (which could have been located in more expensive cities) only once they were certain they would receive the financial aid.

Persistence. Third, I investigate the effect of eligibility to the aide au mérite in the Bac year on persistence in higher education. Specifically, I assess whether eligibility had an effect on (i) being enrolled in 2nd year in Bac year + 1, (ii) being enrolled in 3rd year in Bac year + 2, (iii) the total number of years enrolled in higher education, and (iv) the highest level of study attained.²⁸ The latter two outcomes are measured up to 2020, meaning the cohorts are followed for at least 6 years. As in the U.S., persistence in higher education is a particularly important outcome in the French context since dropout rates are relatively high, especially for students in public universities. Only 30% of students who obtained their Bac in 2016 and enrolled in a 3-year bachelors degree (*Licence*) at a public university in 2016 graduated on time in 2019. For students scoring at or above 16/20 at the Bac, this likelihood is greater, at 69%, implying that 30% of high-achieving students fall behind at least one year (Ménard, 2021). Moreover, persistence serves as a useful measurable intermediate outcome for students enrolled in programs for which it is not obvious to measure graduation, typically students enrolled in academic preparatory classes which do not deliver diplomas.

Figures 3.1c-3.1f graphically display the relationship between the various measures of persistence and Bac grade. For all outcomes, this relationship is quite linear and increasing. Enrollment in 2nd year around the 16 threshold is high, at roughly 75%, and it is about 60% for enrollment in 3rd year. These students are, on average, enrolled in higher education for 5.5 years and their highest level of study attained is between 4th and 5th year, corresponding to a master's level degree. As for enrollment, there is no clear discontinuity at 16 for any outcome. Table 3.1 Panel B reports the associated regression discontinuity estimates. Confirming the visual evidence, the donut estimates are close to zero and insignificant. Even the full sample estimates do not point towards large effects, except for enrollment in 3rd year as the unadjusted students perform particularly poorly on this outcome. Overall, these results suggest that being eligible to the aide au mérite in the Bac year has no effect on future persistence in higher education.

²⁸I also report results for *any* enrollment in 2nd and 3rd year in Appendix Tables C.4 and C.6. The estimates are qualitatively similar to those presented in Table 3.1.

Graduation. Fourth, I assess whether eligibility to the aide au mérite in the Bac year had an impact on degree completion. I focus on degree completion at any point in time (between 2009 and 2019) since depending on the chosen degree, the time to graduation will differ. Appendix Figure A.8 graphically displays the results. The data are significantly more noisy and unexpectedly the likelihood of obtaining a degree is decreasing in Bac grade above 16, likely reflecting the imperfect coverage of the graduation data. The regression discontinuity estimates from the donut specification in Table 3.1's Panel C suggest there may be a small negative effect of eligibility though the coefficient's significance depends on the chosen bandwidth. This is consistent with the results on persistence in 3rd year which act as a good proxy for graduation.


Academic Performance. Lastly, I analyse whether eligibility to the aide au mérite in Bac year may have affected the academic performance of students during their studies. Indeed, it may well be that students obtaining around 16/20 at the Bac are sufficiently good academically to be able to pass on to the next academic year regardless of their financial situation. However, having to work while studying may affect students' grades. Since the administrative data does not contain any information on how well students performed²⁹, I investigate this hypothesis by assessing the effect on (i) enrollment in any masters degrees (defined as degrees for which the final year of study is 4 or 5), (ii) enrollment in a selective masters degrees (defined as masters degrees from engineering, business and other private schools), and (iii) the quality of the masters degree (defined in the same way as for undergraduate degrees). The reasoning is that selective or high quality masters degrees choose students mostly based on their undergraduate grades. Both the visual evidence of Figures 3.2a-3.2c and the estimates reported in Table 3.1's Panel D suggest eligibility to the aide au mérite in the Bac year had no effect on these proxies for academic performance.

5.2. *Non-Academic Outcomes: Geographic Mobility*

As we have seen, eligibility to the aide au mérite in the Bac year appears to not have had any effect on various higher education outcomes, such as enrollment, degree quality, persistence, graduation or academic performance. Another margin that the aide au mérite may have impacted is geographic mobility. Indeed evidence from U.S. state-based financial aid programs suggest they are effective in keeping students in-state, thus affect-

²⁹The only measure of academic performance in the data is how many ECTS credits students obtained in the past year, though it is only available for students at public universities.

Figure 3.2. Higher Education Outcomes as a Function of Bac Grade Around the Aide au Mérite Eligibility Threshold (16/20)



Notes: This figure shows non-parametric binned scatter plots of the relationship between various higher education outcomes and Bac grade, around the aide au mérite eligibility threshold (16/20 Bac grade). The higher education outcomes are reported in the subfigure captions. The sample used corresponds to students from the 2009-2014 Bac cohorts who obtained the Bac, filed a financial aid application and were eligible to a need-based grant. Each red bubble corresponds to the average outcome value for students with Bac grade in $[X, X + 0.05)$, with the size of the bubble corresponding to the number of observations in that grade range. The grey bubbles represent the excluded donut observations, that is observations in $[15.7, 16.05]$. The black fitted lines correspond to local linear regressions with a triangular kernel on each side of the threshold, using the (15, 17) bandwidth, and excluding donut observations. Note that each subfigure has its own y-axis scale.

ing eligible students' geographic mobility (Cohodes and Goodman, 2014; Sjoquist and Winters, 2015; Fitzpatrick and Jones, 2016; Bettinger et al., 2019).

As tuition fees are largely inexpensive in the French higher education system (and exempted from for need-based grant eligible students), living costs represent by far the biggest financial burden imposed on students and their families. A large share of this financial burden is captured by housing costs through rents in the case where students do not live with their parents. The aide au mérite may enable students to study further from home and in particular in large cities where rents are high such as Paris, Marseille or Lyon, France's three largest cities. These cities also tend to concentrate many high-quality higher education institutions.

I investigate this in Table 3.3 by assessing whether eligibility to the aide au mérite had an effect on enrolling in a higher education institution located in (i) Paris (Urban Unit), (ii) Paris, Marseille or Lyon (Urban Units), and (iii) a city with over 200,000 inhabitants³⁰. I

³⁰These cities are: Paris, Marseille, Lyon, Toulouse, Nice, Nantes, Montpellier, Strasbourg, Bordeaux, Lille, and Rennes.

assess geographic mobility both in the Bac year and in the following year as students may decide to change cities only once they are certain that they receive the aide au mérite. The estimates suggest eligibility to the aide au mérite may have had slightly positive effects on geographic mobility towards big cities. The effects for the year following the Bac year are of the same magnitude, implying all the potential effect on geographic mobility is driven by location decisions made in the Bac year, not once students were certain they would receive the aide au mérite. The visual evidence presented in Figure A.9 does seem to point towards some slight effects on geographic mobility though it is not striking either.

Table 3.3: Effects of Eligibility to the Aide au Mérite in Bac Year on Location of Studies

	Mean [15.5, 15.7]	Point Estimate	95% Confidence Interval	Bandwidth	# obs. left	# obs. right
	(1)	(2)	(3)	(4)	(5)	(6)
<i>Panel A. In Bac Year</i>						
Paris (Urban Unit)	0.13	0.03*	[-0.01, 0.07]	(15.35, 16.65)	15,810	18,367
Paris (Urban Unit)	0.13	0.02**	[0, 0.04]	(15, 17)	35,234	25,357
Paris, Marseille or Lyon (Urban Units)	0.21	0.04**	[0, 0.09]	(15.31, 16.69)	17,931	19,136
Paris, Marseille or Lyon (Urban Units)	0.21	0.028***	[0.01, 0.05]	(15, 17)	35,234	25,357
City Over 200k Inhabitatnts	0.35	0.033	[-0.01, 0.07]	(15.17, 16.83)	25,535	22,272
City Over 200k Inhabitatnts	0.35	0.029**	[0, 0.05]	(15, 17)	35,234	25,357
<i>Panel B. In Bac Year + 1</i>						
Paris (Urban Unit)	0.13	0.04*	[-0.01, 0.09]	(15.39, 16.61)	13,831	17,209
Paris (Urban Unit)	0.13	0.023***	[0.01, 0.04]	(15, 17)	35,234	25,357
Paris, Marseille or Lyon (Urban Units)	0.21	0.028	[-0.01, 0.08]	(15.3, 16.7)	17,931	19,136
Paris, Marseille or Lyon (Urban Units)	0.21	0.023**	[0, 0.04]	(15, 17)	35,234	25,357
City Over 200k Inhabitatnts	0.34	0.059**	[0.01, 0.11]	(15.27, 16.73)	19,442	19,862
City Over 200k Inhabitatnts	0.34	0.038***	[0.01, 0.06]	(15, 17)	35,234	25,357

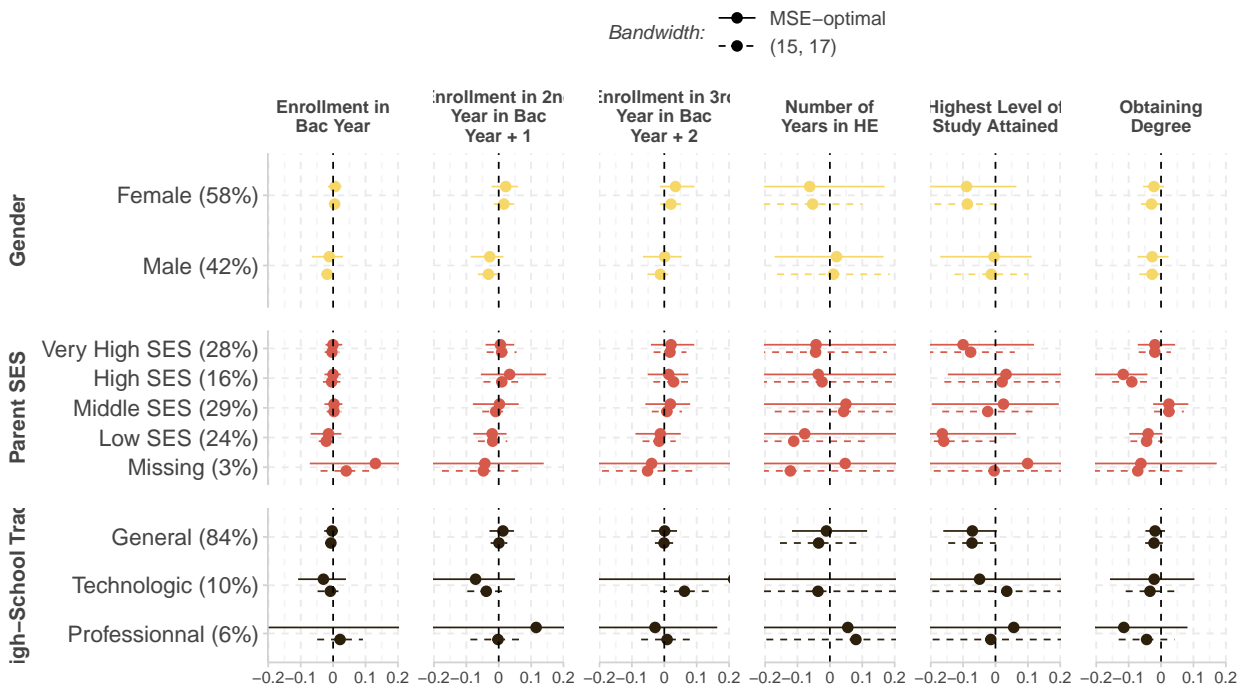
Notes: This table reports estimates of the discontinuity in the location of studies, in Bac year (panel A) and in the following year (panel B), at the aide au mérite eligibility threshold (16/20 Bac grade), excluding students with Bac grades in [15.7, 16.05]. The higher education institution location are reported in the first column's rows. The cities with over 200,000 inhabitants are: Paris, Marseille, Lyon, Toulouse, Nice, Nantes, Montpellier, Strasbourg, Bordeaux, Lille, and Rennes. Column (1) reports the mean outcome for students with Bac grade in [15.5, 15.7]. Estimates (col. (2)) are reported for both the MSE-optimal (upper row) and the (15, 17) (lower row) bandwidths. The MSE-optimal bandwidth, obtained using the *rdrobust* R package, varies across each outcome. For MSE-optimal bandwidth estimates, the associated confidence intervals (col. (3)) correspond to robust 95% confidence intervals, while they correspond to conventional 95% confidence intervals for (15, 17) bandwidth estimates. Column (4) reports the bandwidth over which the local linear regressions are estimated (it is centered around 16), while columns (5) and (6) report, respectively, the number of observations used for estimation to the left and right of the 16/20 threshold. See Table 2.1's notes for an example of how to read the estimates. Statistical significance is computed based on the relevant p-value and ***, **, and * indicate significance at 1, 5, and 10%, respectively. All the detailed regression tables can be found in Appendix C.

5.3. Heterogeneity

Before explicitly attempting to uncover potential mechanisms, I explore the heterogeneity of the effects by various characteristics such as gender, SES and high school track. The objective is in a sense to find out whether a particular subgroup of students might have been more responsive to the additional financial aid awarded by the aide au mérite. Fig-

Figure 3.3 displays the point estimates and associated 95% confidence intervals for the main outcomes, for both the MSE-optimal (solid line) and (15, 17) bandwidths (dashed line). The results indicate eligibility to the aide au mérite had no statistically differential effect between men and women, with the coefficients being small in magnitude and statistically insignificant. There is also no clear differential effect across SES, with small point estimates (though the confidence intervals are sometimes large). The same applies for high school track, though sample sizes for the technologic and professional tracks are small. This points towards the null effects found previously reflecting true nulls rather than hidden heterogeneous effects averaging out.

Figure 3.3. Heterogeneity of Effects of Eligibility to the Aide au Mérite in Bac Year by Gender, SES, and High School Track



Notes: This figure shows estimates of the discontinuity in higher education outcomes at the aide au mérite eligibility threshold (16/20 Bac grade), for subsamples of students, excluding students with Bac grades in [15.7, 16.05]. The higher education outcomes are reported in column's titles, while the subsamples (and % in the full sample) are reported in the first column's rows. The horizontal axis is in percentage points. Estimates for two different bandwidths are reported: the MSE-optimal and (15, 17) bandwidths. The MSE-optimal bandwidth, obtained using the *rdrobust* R package, varies across each outcome. For MSE-optimal bandwidth estimates, associated confidence intervals correspond to robust 95% confidence intervals, while they correspond to conventional 95% confidence intervals for (15, 17) bandwidth estimates. For example, the first row of estimates displays the discontinuity in each higher education outcome at the 16/20 Bac grade threshold on the subsample of students who are female (who represent 58% of the full sample).

6. Mechanisms

Several mechanisms may explain the null results found in the main analysis. First, the null effects on enrollment could be due to students being unaware that they would be receiving this additional financial aid. As such no behavioral response could take place. Second, the additional aid may have crowded out students' parents' financial assistance, in which case on net eligible students were not better positioned financially compared to barely non-eligible students. Third, the amount of the aide au mérite may have been too small relative to need-based grants, therefore not inducing behavioral responses beyond those from eligibility to need-based grants. Lastly, high-achieving, low-income students may be less sensitive to financial aid, since they likely have high expected returns to education and intrinsic motivation, and therefore strong reasons to attend and persist in higher education.

6.1. *Were Students Unaware of the Policy?*

Since the aide au mérite was introduced at the same time as a vast reform of the need-based grant system it might have attracted less attention and therefore been less salient to eligible students. To test whether information about eligibility to the program might have played a role, I assess whether students who may have been better informed exhibited larger behavioral responses. In particular, I estimate whether effects were larger (i) for more recent Bac cohorts, who are reasonably expected to have been better informed, (ii) for students who attended a high school which had more more students in the previous Bac cohort who received the aide au mérite³¹, and (iii) for students who had more high school peers who were also eligible to the aide au mérite in the same year.

Table 3.1 displays the results for these three tests. The results show no difference in the likelihood of enrollment across all proxies for information (all estimates are very close to zero and overwhelmingly not significant), strongly suggesting lack information was likely not the main driver of the null effects of eligibility to the aide au mérite on enrollment.

6.2. *Did the Aid Crowd Out Parents' Financial Support?*

Another potential explanation for the null results is that the aide au mérite crowded out parents' financial support. From a policy perspective, these behavioral responses from

³¹By construction, this test excludes the 2009 Bac cohort.

Table 3.1: Effects of Eligibility to the Aide au Mérite on Enrollment in Bac Year by Student Awareness Proxy

	Mean [15.5, 15.7)	<i>Bandwidth:</i>			
		MSE-Optimal		(15, 17)	
		(1)	(2)	(3)	(4)
<i>Panel A: Bac Cohort</i>					
2009	0.94	0.003 [-0.03, 0.03]	0.002 [-0.03, 0.04]	-0.001 [-0.03, 0.03]	0.001 [-0.03, 0.03]
2010	0.94	0.004 [-0.05, 0.05]	0.003 [-0.05, 0.05]	0.002 [-0.03, 0.04]	0.002 [-0.03, 0.03]
2011	0.94	-0.004 [-0.05, 0.04]	0.003 [-0.04, 0.04]	0.002 [-0.03, 0.03]	0.006 [-0.02, 0.04]
2012	0.94	-0.007 [-0.05, 0.03]	-0.01 [-0.06, 0.04]	-0.011 [-0.04, 0.02]	-0.014 [-0.04, 0.01]
2013	0.95	0.002 [-0.03, 0.03]	0.003 [-0.02, 0.03]	0.002 [-0.02, 0.03]	0.003 [-0.02, 0.03]
2014	0.95	-0.02 [-0.06, 0.01]	-0.017 [-0.05, 0.01]	-0.018 [-0.04, 0.01]	-0.019 [-0.04, 0.01]
<i>Panel B: Number of Recipients in Same High School in Previous Cohort</i>					
0-2 (33.7%)	0.92	0.01 [-0.01, 0.03]	0.009 [-0.01, 0.03]	0.01 [-0.02, 0.04]	0.009 [-0.02, 0.03]
3-5 (28.7%)	0.96	-0.009 [-0.04, 0.02]	-0.005 [-0.03, 0.02]	-0.009 [-0.03, 0.01]	-0.009 [-0.03, 0.01]
6-9 (22.3%)	0.96	0.02 [-0.03, 0.06]	0.019 [-0.03, 0.07]	-0.003 [-0.03, 0.02]	-0.003 [-0.03, 0.02]
10+ (15.2%)	0.96	-0.03* [-0.07, 0]	-0.028 [-0.07, 0.01]	-0.029** [-0.06, 0]	-0.029** [-0.06, 0]
<i>Panel C: Number of Other Eligible Students in Same High School in Same Cohort</i>					
0-2 (27.3%)	0.95	-0.044** [-0.09, -0.01]	-0.025 [-0.07, 0.01]	-0.037*** [-0.06, -0.01]	-0.018 [-0.04, 0.01]
3-5 (27.1%)	0.96	-0.001 [-0.03, 0.03]	-0.003 [-0.03, 0.02]	-0.001 [-0.02, 0.02]	-0.003 [-0.02, 0.02]
6-9 (23.7%)	0.95	-0.002 [-0.05, 0.04]	0 [-0.05, 0.04]	-0.005 [-0.03, 0.02]	-0.005 [-0.03, 0.02]
10+ (21.8%)	0.95	-0.006 [-0.03, 0.01]	-0.004 [-0.03, 0.02]	0.002 [-0.02, 0.03]	0.003 [-0.02, 0.03]
Controls			✓		✓

Notes: This table reports estimates of the discontinuity in enrollment in Bac year at the aide au mérite eligibility threshold (16/20 Bac grade), for subsamples of students, excluding students with Bac grades in [15.7, 16.05]. Panel A reports estimates for each Bac cohort separately. Panel B reports estimates by students' number of aide au mérite recipients from the same high school in the previous Bac cohort. Panel C reports estimates by students' number of aide au mérite eligibles in the same high school in the same Bac cohort. Estimates for two different bandwidths are reported: the MSE-optimal and (15, 17) bandwidths. The MSE-optimal bandwidth, obtained using the *rdrobust* R package, varies across each outcome. For MSE-optimal bandwidth estimates, the ranges in brackets correspond to associated robust 95% confidence intervals, while they correspond to associated conventional 95% confidence intervals for (15, 17) bandwidth estimates. Control variables are gender, age, SES, Bac track, and Bac cohort. Column (1) reports the mean of the row's subsample enrollment in Bac year for students with Bac grade in [15.5, 15.7). See Table 3.2's notes for an example of how to read the estimates. Statistical significance is computed based on the relevant p-value and ***, **, and * indicate significance at 1, 5, and 10%, respectively.

parents should be taken into account. However, it is important to disentangle whether this particular behavioral response is the main explanation for the results found. In particular, could it be that students who were eligible to the aide au mérite and those who were not ended up equally well off financially as the former's parents reduced their financial contributions by exactly 200 euros per month.

Since I cannot directly observe parents' financial assistance, to investigate this channel I analyse whether the most financially disadvantaged students experienced benefits from eligibility to the aide au mérite. The idea being that such students were likely to receive little or no financial support from their families and therefore even full crowding out could not completely compensate the extra 200 euros. In particular, I assess the effects of eligibility to the aide au mérite for students in the bottom 10% of the parent income distribution as well as for echelons 6 and 7 students³². Using survey data from 2014, [Grobon and Wolff \(2022\)](#) find that echelon 7 students, the most disadvantaged, received (on average) roughly 100 euros from their parents each month, while echelon 6 students received about 150 euros. Thus even with full parent crowding out, these students would still, on average, receive a net financial gain from the aide au mérite.

Table 3.2 displays the results from this analysis. The estimates for both students in the bottom 10% of the parent income distribution and for echelon 6 and 7 students are small in magnitude and insignificant across all outcomes. That being said, the confidence intervals are non-trivial due to the small sample size. Nonetheless, these results are indicative that crowding out of parents' financial assistance is unlikely to have been the main driver of the overall null effects. Of course, this test cannot rule out interactions between parent income and eligibility to the aide au mérite other than through crowding out.

6.3. *Was the Awarded Amount Too Small Relative to Need-Based Grants?*

One may be concerned that the amount being awarded by the aide au mérite was too small relative to the amount of need-based grants to induce behavioral responses. Eligible students other than those at echelon 0 of the need-based grant scheme were already eligible to some financial aid (see Appendix Table B.2), thus perhaps muting any potential effect of the aide au mérite. Theoretically, the mechanism underlying this line of thinking is that above a certain amount financial aid has no marginal effect on students' behavior.

I test for this hypothesis by analysing whether students at echelon 0, 0 bis and 1 exhibited differential effects from eligibility to the aide au mérite. For echelon 0 students,

³²Due to sample size restrictions it is not possible to estimate the effects only for echelon 7 students.

Table 3.2: Test of Crowding Out Effects of Aide au Mérite Eligibility

	Enrollment in Bac Year	Enrollment in 2nd Year in Bac Year + 1	Enrollment in 3rd Year in Bac Year + 2	Number of Years in Higher Education	Highest Level of Study Attained	Obtaining a Degree
	(1)	(2)	(3)	(4)	(5)	(6)
<i>Panel A. Parent Income in Bottom 10%</i>						
Mean [15.5, 15.7)	0.92	0.66	0.47	4.9	3.85	0.54
MSE-Optimal	0.024	0.02	-0.01	0.412	0.441	0.019
(15, 17)	[-0.03, 0.09]	[-0.12, 0.17]	[-0.13, 0.09]	[-0.37, 1.24]	[-0.13, 1.16]	[-0.15, 0.19]
	0.022	0.007	0.014	0.345	0.336*	-0.013
	[-0.03, 0.08]	[-0.09, 0.11]	[-0.1, 0.12]	[-0.19, 0.88]	[-0.02, 0.7]	[-0.12, 0.1]
<i>Panel B. Echelons 6 & 7</i>						
Mean [15.5, 15.7)	0.93	0.69	0.49	5.01	3.98	0.6
MSE-Optimal	0.006	-0.001	0.002	0.199	0.102	-0.056
(15, 17)	[-0.03, 0.04]	[-0.09, 0.07]	[-0.07, 0.07]	[-0.32, 0.72]	[-0.22, 0.42]	[-0.17, 0.03]
	0.003	-0.001	0.022	0.189	0.095	-0.064*
	[-0.04, 0.04]	[-0.07, 0.07]	[-0.05, 0.1]	[-0.17, 0.55]	[-0.15, 0.34]	[-0.14, 0.01]

Notes: This table reports estimates of the discontinuity in higher education outcomes at the aide au mérite eligibility threshold (16/20 Bac grade), excluding students with Bac grades in [15.7, 16.05]. Panel A reports estimates on the subsample of students with parents in the bottom 10% of the parent income distribution in Bac year. Panel B reports estimates on the subsample of students with need-based grant echelons 6 or 7 in Bac year. For each panel, the first row reports the average outcome for students with Bac grades in [15.5, 15.7), the second row reports estimates obtained using the MSE-optimal bandwidth (using the *rdrobust* R package), and the third row reports estimates obtained using the (15, 17) bandwidth. For MSE-optimal bandwidth estimates, the ranges in brackets correspond to associated robust 95% confidence intervals, while they correspond to associated conventional 95% confidence intervals for (15, 17) bandwidth estimates. For example, for students in the bottom 10% of the parent income distribution in Bac year, column (1) indicates that, using the MSE-optimal bandwidth (without controls), the estimated discontinuity in enrollment in Bac year around 16/20 is 2.4 percentage points, with the share of such students enrolled among students with Bac grade in [15.5, 15.7) being 92%. This discontinuity estimate is 2.2 percentage points when using the (15, 17) bandwidth. Statistical significance is computed based on the relevant p-value and ***, **, and * indicate significance at 1, 5, and 10%, respectively.

non-eligible students only benefit from tuition fee exemption while eligible students also receive 200 euros per month. Echelon 0 bis students receive 1,000 euros annually, while echelon 1 students receive roughly 1,500 euros, i.e., about the same amount as the aide au mérite (1,800 euros). Table 3.3 presents the results from this analysis. Once more, all the estimates are very small in magnitude (except for enrollment in 3rd year for echelon 1 students) for students at these three echelons. There is therefore no difference in effects between students receiving no cash allowance as part of their need-based grant (echelon 0), those receiving small amounts (echelon 0 bis) and those receiving an amount roughly equal to the amount of the aide au mérite (echelon 1). These results hold when restricting to students whose parents' incomes are within 5% of the echelon 1 threshold in 2009-2012 (before the introduction of the 0 bis echelon), as shown in Appendix Table B.7, though the confidence intervals are large due to small sample size. This implies that the fact that the aide au mérite was awarded on top of existing financial aid is probably not the main explanation for the null effects.

Table 3.3: Test of Whether aide au mérite Eligibility Had Effects for Students with No or Small Amounts of Cash Allowance as Part of their Need-Based Grant

	Enrollment in Bac Year	Enrollment in 2nd Year in Bac Year + 1	Enrollment in 3rd Year in Bac Year + 2	Number of Years in Higher Education	Highest Level of Study Attained	Obtaining a Degree
	(1)	(2)	(3)	(4)	(5)	(6)
<i>Panel A. Echelon 0 (No Cash Allowance)</i>						
Mean [15.5, 15.7)	0.95	0.74	0.61	5.53	4.52	0.64
MSE-Optimal	-0.002	-0.011	0.007	0.005	-0.042	-0.006
(15, 17)	[-0.02, 0.02]	[-0.1, 0.06]	[-0.07, 0.07]	[-0.24, 0.24]	[-0.25, 0.16]	[-0.09, 0.08]
	-0.01	-0.016	-0.001	0.028	-0.048	-0.009
	[-0.03, 0.01]	[-0.06, 0.03]	[-0.05, 0.05]	[-0.2, 0.25]	[-0.19, 0.09]	[-0.06, 0.04]
<i>Panel B. Echelons 0 bis (1,000 Euros of Annual Cash Allowance)</i>						
Mean [15.5, 15.7)	0.95	0.74	0.61	4.94	4.28	0.61
MSE-Optimal	0.002	-0.026	-0.022	0.088	0.012	-0.012
(15, 17)	[-0.06, 0.06]	[-0.09, 0.03]	[-0.16, 0.12]	[-0.21, 0.4]	[-0.21, 0.22]	[-0.09, 0.09]
	0	-0.008	-0.01	0.077	0.038	-0.015
	[-0.04, 0.04]	[-0.08, 0.07]	[-0.09, 0.07]	[-0.24, 0.39]	[-0.21, 0.29]	[-0.1, 0.07]
<i>Panel C. Echelons 1 (1,500 Euros of Annual Cash Allowance)</i>						
Mean [15.5, 15.7)	0.94	0.73	0.56	5.27	4.32	0.62
MSE-Optimal	0.003	0.027	0.111***	0.211	0.22	-0.014
(15, 17)	[-0.02, 0.03]	[-0.03, 0.09]	[0.04, 0.21]	[-0.37, 0.86]	[-0.08, 0.56]	[-0.08, 0.06]
	0.014	0.028	0.081***	0.062	0.042	-0.013
	[-0.01, 0.04]	[-0.02, 0.08]	[0.02, 0.14]	[-0.2, 0.32]	[-0.13, 0.21]	[-0.07, 0.04]

Notes: This table reports estimates of the discontinuity in higher education outcomes at the aide au mérite eligibility threshold (16/20 Bac grade), excluding students with Bac grades in [15.7, 16.05]. Panel A reports estimates on the subsample of students with need-based grant echelon 0, i.e., who received no cash allowance. Panel B reports estimates on the subsample of students with need-based grant echelon 0 bis, i.e., who received a cash allowance of roughly 1,000 euros annually. Panel C reports estimates on the subsample of students with need-based grant echelon 1, i.e., who received a cash allowance of roughly 1,500 euros annually. For each panel, the first row reports the average outcome for students with Bac grades in [15.5, 15.7), the second row reports estimates obtained using the MSE-optimal bandwidth (using the *rdrobust* R package), and the third row reports estimates obtained using the (15, 17) bandwidth. For MSE-optimal bandwidth estimates, the ranges in brackets correspond to associated robust 95% confidence intervals, while they correspond to associated conventional 95% confidence intervals for (15, 17) bandwidth estimates. See Table 3.2's notes for an example of how to read the estimates. Statistical significance is computed based on the relevant p-value and ***, **, and * indicate significance at 1, 5, and 10%, respectively.

6.4. Discussion: Were Targeted Students Not Marginal Students?

The last plausible explanation - given that the lack of information, crowding out of parent financial assistance and the low amount of the aide au mérite relative to need-based grants do not appear to explain the null effects - is simply that targeted students were not marginal students, in the sense of students whose enrollment or persistence behavior rests on eligibility to financial aid. Indeed, the aide au mérite targeted students in top 5% of the Bac grade distribution. Such students are very able academically and are likely set on pursuing higher education largely regardless of the available financial assistance they are eligible to since their expected returns to higher education are large and they are academically inclined.

This is consistent with several studies that examine heterogeneity of the effects of financial aid by the academic ability of eligible students. Goodman (2008) was the first to observe that the effects of financial aid (in his case, Massachusetts's Adams Scholarship) varied by

Table 3.4: Summary of Studies Finding Complementarities Between Financial Aid and Academic Ability

Study	Program	Outcome	Student Academic level	Effect (p.p.)	Baseline (%)
Cohodes and Goodman (2014, Table 7)	Massachusetts Adams Scholarship (<i>merit-based scholarship covering tuition at in-state public college</i>)	Enrollment	Lower income and less academically skilled	+7.7***	–
			Higher income and more academically skilled	+1.5	–
Fack and Grenet (2015, Table 4)	French need-based grants (<i>need-based grant of 1,500 euros</i>)	Enrollment	Bottom quartile of Bac grade distribution	+3.4***	75
			Top quartile of Bac grade distribution	+1.8*	78.5
Bettinger et al. (2019, Tabel 2)	California Cal Grant (<i>merit/need-based grant covering 4 years of tuition assistance for in-state HEI</i>)	BA completion	Students around the GPA cutoff (GPA = 3.08/4)	+4.6*** p.p.	46
			Students around the income cutoff (GPA = 3.55/4)	+3	67
Angrist et al. (2022, Figure 4)	Nebraska STBF Scholarship (<i>merit/need-based scholarship covering college costs at in-state public college</i>)	BA completion	Students with below-median GPA	+12***	42
			Students with above-median GPA	+4**	80

skill level, and highlighted its importance in explaining varying effects of financial aid.

In Table 3.4, I compile the results from all the quasi-experimental studies that, to my knowledge, examine such heterogeneity by academic level. Across very different states or countries, types of financial aid (need-based, merit-based or both) and outcomes (enrollment or BA graduation), this (allegedly small) set of studies consistently finds that higher-achieving students benefit less from financial aid than their lower achieving counterparts. The only exception I have found is Castleman and Long (2016), who find significantly larger effects of eligibility to Florida’s Student Access Grant (FSAG) on degree attainment for students with higher high school GPAs, though these students also received slightly larger FSAG amounts, complicating the interpretation of the results slightly.

My findings are consistent with these results, and thus point to the fact that there may be complementarities between financial aid and academic ability. In other words, students with lower academic levels at the point of college entry are likely to be significantly more adversely impacted by a lack of financial support relative to students with greater college readiness. This is suggestive that financial aid targeted at lower-achieving students may yield greater effect, though one should be wary of the perverse consequences such a system may have if it leads high-achieving students to perform poorly for financial aid eligibility reasons.

7. Conclusion

This paper examines whether automatically providing additional financial aid to high achieving, low-income students can reduce the higher education enrollment and degree

quality gap with their high-income peers. I find that the aide au mérite, a financial aid scheme introduced in France that gave an additional 1,800 euros annually to low-income students scoring in the top 5% of the national end of high school exam and who enrolled in higher education, appears to have had no discernible effect on enrollment, degree quality, persistence, graduation nor academic performance in higher education.

In trying to uncover what may be the potential mechanisms explaining these null effects, I find evidence suggesting they were not driven by students being unaware of their eligibility, crowding out of parents' financial support nor that the amount awarded was too small relative to other financial aid that eligible students also received. The most plausible explanation is that high-achieving students, even from low-income backgrounds, are not marginal students in the sense that they are not indifferent between enrolling in higher education or not absent financial aid.

By putting these results in light of heterogeneity analyses across a number of studies, I find consistent evidence pointing towards complementarities between financial aid and college "unpreparedness", with students with lower academic levels at the point of college entry being likely to be significantly more adversely impacted by a lack of financial support relative to students with greater college readiness. This potentially suggests that merit aid targeting very high-achieving students may perhaps be inadequate in improving their higher education outcomes, and this aid could likely benefit more students lower down the ability distribution.

It should be highlighted that financial aid not only affects academic, geographic or socio-economic outcomes but may very well impact unobserved outcomes such as students' mental health and financial distress. This is an outcome which has received very little attention in the literature due to data limitations but may be of great importance nonetheless.

References

- Andrews, Rodney J., Scott A. Imberman, and Michael F. Lovenheim**, “Recruiting and Supporting Low-Income, High-Achieving Students at Flagship Universities,” *Economics of Education Review*, February 2020, 74, 101923.
- Angrist, Joshua D, Victor Lavy, Jetson Leder-Luis, and Adi Shany**, “Maimonides Rule Redux,” *American Economic Review: Insights*, 2019, 1 (3), 16.
- Angrist, Joshua, David Autor, and Amanda Pallais**, “Marginal Effects of Merit Aid for Low-Income Students,” *The Quarterly Journal of Economics*, April 2022, 137 (2), 1039–1090.
- Barr, Andrew, Jonathan Eggleston, and Alexander A Smith**, “Investing in Infants: The Lasting Effects of Cash Transfers to New Families*,” *The Quarterly Journal of Economics*, November 2022, 137 (4), 2539–2583.
- Barreca, Alan I., Jason M. Lindo, and Glen R. Waddell**, “Heaping-Induced Bias in Regression-Discontinuity Designs,” *Economic Inquiry*, 2016, 54 (1), 268–293.
- Baumgartner, Hans J. and Viktor Steiner**, “Does More Generous Student Aid Increase Enrolment Rates into Higher Education? Evaluating the German Student Aid Reform of 2001,” *SSRN Electronic Journal*, 2006.
- Bechichi, Nagui and Georgia Thebault**, “Students’ Preferences, Capacity Constraints and Post-Secondary Achievements in a Non-Selective System,” Technical Report N° G2021/01, INSEE, Paris January 2021.
- Bettinger, Eric, Oded Gurantz, Laura Kawano, Bruce Sacerdote, and Michael Stevens**, “The Long-Run Impacts of Financial Aid: Evidence from California’s Cal Grant,” *American Economic Journal: Economic Policy*, 2019, 11 (1), 64–94.
- Bharadwaj, Prashant, Katrine Velleesen Løken, and Christopher Neilson**, “Early Life Health Interventions and Academic Achievement,” *American Economic Review*, August 2013, 103 (5), 1862–1891.
- Black, Sandra E., Jeffrey T. Denning, and Jesse Rothstein**, “Winners and Losers? The Effect of Gaining and Losing Access to Selective Colleges on Education and Labor Market Outcomes,” *American Economic Journal: Applied Economics*, January 2023, 15 (1), 26–67.
- Bleemer, Zachary**, “Affirmative Action, Mismatch, and Economic Mobility after California’s Proposition 209,” *The Quarterly Journal of Economics*, December 2021, 137 (1), 115–160.
- Bonneau, Cécile, Pauline Charousset, Julien Grenet, and Georgia Thebault**, “Quelle démocratisation des grandes écoles depuis le milieu des années 2000?,” Technical Report 30, Institut des Politiques Publiques, Paris 2021.


- Burland, Elizabeth, Susan Dynarski, Katherine Micheltore, Stephanie Owen, and Shwetha Raghuraman**, “The Power of Certainty: Experimental Evidence on the Effective Design of Free Tuition Programs,” *AER: Insights*, 2023, 5 (3), 293–310.
- Calonico, Sebastian, Matias D. Cattaneo, and Rocio Titiunik**, “Robust Nonparametric Confidence Intervals for Regression-Discontinuity Designs: Robust Nonparametric Confidence Intervals,” *Econometrica*, November 2014, 82 (6), 2295–2326.
- , – , **Max H. Farrell, and Rocío Titiunik**, “Regression Discontinuity Designs Using Covariates,” *The Review of Economics and Statistics*, July 2019, 101 (3), 442–451.
- Campbell, Stuart, Lindsey Macmillan, Richard Murphy, and Gill Wyness**, “Matching in the Dark? Inequalities in Student to Degree Match,” *Journal of Labor Economics*, October 2022, 40 (4), 807–850.
- Canaan, Serena and Pierre Mouganie**, “Returns to Education Quality for Low-Skilled Students: Evidence from a Discontinuity,” *Journal of Labor Economics*, 2018, 36 (2), 395–436.
- Carrell, Scott and Bruce Sacerdote**, “Why Do College-Going Interventions Work?,” *American Economic Journal: Applied Economics*, July 2017, 9 (3), 124–151.
- Castleman, Benjamin L and Bridget Terry Long**, “Looking beyond Enrollment: The Causal Effect of Need-Based Grants on College Access, Persistence, and Graduation,” *Journal of Labor Economics*, 2016, 34 (4), 52.
- Cattaneo, Matias D., Nicolás Idrobo, and Rocío Titiunik**, *A Practical Introduction to Regression Discontinuity Designs: Foundations* Cambridge Elements: Quantitative and Computational Methods for Social Science, 1 ed., Cambridge University Press, November 2019.
- Chetty, Raj, David J Deming, and John N Friedman**, “Diversifying Society’s Leaders? The Determinants and Causal Effects of Admission to Highly Selective Private Colleges,” NBER Working Paper 31492, National Bureau of Economic Research 2023.
- Cohodes, Sarah R. and Joshua S. Goodman**, “Merit Aid, College Quality, and College Completion: Massachusetts’ Adams Scholarship as an In-Kind Subsidy,” *American Economic Journal: Applied Economics*, October 2014, 6 (4), 251–285.
- Crawford, Claire, Paul Gregg, Lindsey Macmillan, Anna Vignoles, and Gill Wyness**, “Higher Education, Career Opportunities, and Intergenerational Inequality,” *Oxford Review of Economic Policy*, 2016, 32 (4), 553–575.
- DARES**, “Les Activités Rémunérées Des Étudiants : Quelles Formes et Quelle Organisation ?,” DARES Analyses 046 2017.
- Dearden, Lorraine, Emla Fitzsimons, and Gill Wyness**, “Money for Nothing: Estimating the Impact of Student Aid on Participation in Higher Education,” *Economics of Education Review*, December 2014, 43, 66–78.

- Deshpande, Manasi, Tal Gross, and Yalun Su**, “Disability and Distress: The Effect of Disability Programs on Financial Outcomes,” *American Economic Journal: Applied Economics*, April 2021, 13 (2), 151–178.
- Dynarski, Susan, Cj Libassi, Katherine Micheltore, and Stephanie Owen**, “Closing the Gap: The Effect of Reducing Complexity and Uncertainty in College Pricing on the Choices of Low-Income Students,” *American Economic Review*, June 2021, 111 (6), 1721–1756.
- Fack, Gabrielle and Julien Grenet**, “Improving College Access and Success for Low-Income Students: Evidence from a Large Need-Based Grant Program,” *American Economic Journal: Applied Economics*, April 2015, 7 (2), 1–34.
- Fitzpatrick, Maria D. and Damon Jones**, “Post-Baccalaureate Migration and Merit-Based Scholarships,” *Economics of Education Review*, October 2016, 54, 155–172.
- Goodman, Joshua**, “Who Merits Financial Aid?: Massachusetts’ Adams Scholarship,” *Journal of Public Economics*, October 2008, 92 (10-11), 2121–2131.
- Grobon, Sébastien and François-Charles Wolff**, “Do Public Scholarships Crowd Out Parental Transfers ? Evidence from France,” *SSRN Electronic Journal*, 2022.
- Hakimov, Rustamdjan, Renke Schmacker, and Camille Terrier**, “Confidence and College Applications: Evidence from a Randomized Intervention,” 2022.
- Hansen, Benjamin**, “Punishment and Deterrence: Evidence from Drunk Driving,” *American Economic Review*, April 2015, 105 (4), 1581–1617.
- Herbaut, Estelle and Koen Geven**, “What Works to Reduce Inequalities in Higher Education? A Systematic Review of the (Quasi-)Experimental Literature on Outreach and Financial Aid,” *Research in Social Stratification and Mobility*, February 2020, 65, 100442.
- Hoxby, Caroline and Christopher Avery**, “The Missing “One-Offs”: The Hidden Supply of High-Achieving, Low-Income Students,” *Brookings Papers on Economic Activity*, 2013, (Spring).
- Hoxby, Caroline M and Sarah Turner**, “Informing Students about Their College Options: A Proposal for Broadening the Expanding College Opportunities Project,” Technical Report 2013-03, The Hamilton Project 2013.
- Imbens, Guido W. and Thomas Lemieux**, “Regression Discontinuity Designs: A Guide to Practice,” *Journal of Econometrics*, February 2008, 142 (2), 615–635.
- Lee, David S and Thomas Lemieux**, “Regression Discontinuity Designs in Economics,” *Journal of Economic Literature*, June 2010, 48 (2), 281–355.
- Ménard, Boris**, “Parcours et Réussite En Licence : Les Résultats de La Session 2020,” Technical Report N° 24, Ministère de l’Enseignement Supérieur, de la Recherche et de l’Innovation, Paris November 2021.

- MESR**, “L’état de l’Enseignement Supérieur et de La Recherche,” Technical Report, Ministère de l’Enseignement Supérieur et de la Recherche, Paris 2011.
- Montalbán, José**, “Countering Moral Hazard in Higher Education: The Role of Performance Incentives in Need-Based Grants,” *The Economic Journal*, forthcoming.
- Murphy, Richard and Gill Wyness**, “Testing Means-Tested Aid,” *Journal of Labor Economics*, March 2022, p. 719995.
- Sjoquist, David L. and John V. Winters**, “State Merit Aid Programs and College Major: A Focus on STEM,” *Journal of Labor Economics*, October 2015, 33 (4), 973–1006.
- Vergolini, Loris, Nadir Zanini, and Nicola Bazoli**, “Liquidity Constraints and University Participation in Times of Recession. Evidence from a Small-scale Programme,” Technical Report No. 2014-11, Research Institute for the Evaluation of Public Policies, Trento, Italy 2014.


A. Appendix Figures

Figure A.1. Aide au mérite Status Over Time




Notes: This figure displays the evolution over time of aide au mérite status for students in 2009-2014 Bac cohorts who were eligible to a need-based grant in their Bac year, and received or should have received the aide au mérite at least once up to Bac year + 2. The first "Bac year" column shows that among these students, the vast majority received the aide au mérite, as they should have, in their Bac year ("Yes: received"). Approximately 5% fulfilled all the necessary criteria (at least 16/20 at the Bac, eligible to a need-based grant, and enrolled in higher education) yet did not receive it ("NO: should have"). The remaining couple of percentages are students who did not receive it in their Bac year but ended up receiving it or being in the group that should have received it in subsequent years ("NO: other*", see the figure's caption for additional details on this category). Among students who received the aide au mérite in their Bac year, a very large fraction continued to receive it the following years.

Figure A.2. Need-Based Grant Eligibility Over Time




Notes: This figure displays the evolution over time of need-based grant eligibility status over time for students in the full sample. The first “Bac year” column shows that all these students were eligible to a need-based grant (echelons 0 to 7), which is expected since the full sample contains only students eligible to a need-based grant in their Bac year. In the following years, the vast majority who file a financial aid application remain eligible while some do not enrol in higher education and do not file a financial aid application.

Figure A.3. Distribution of Bac grades, 2009-2014, all students




Notes: This figure shows the distribution of Bac grades of all students in 2009-2014 Bac cohorts, in panel (a) between 8 and 20, and in panel (b) between 15 and 17. For panel (a) each bin represents the number of students who obtained a Bac grade in $[X, X + 0.1)$, while in panel (b) each bin represents the number of students who obtained a Bac grade in $[X, X + 0.05)$.

Figure A.4. Distribution of Bac grades, 2009-2014 - 16.05 Grades In [16, 16.05) Bin



Notes: This figure shows the distribution of Bac grades between 15 and 17 for 2009-2014 Bac cohorts. Each bin represents the number of students who obtained a Bac grade in $[X, X + 0.05)$, except for the $[16, 16.05]$ bar which includes students scoring at 16.05.

Figure A.5. Discontinuity Estimates on Observable Characteristics for Different Donut Limits



Notes: This figure shows the estimated discontinuities in students' observable characteristics (on the vertical axis) around the aide au mérite eligibility threshold (16/20 Bac grade) for different donut boundaries.

Figure A.6. Discontinuity in Pre-Treatment Observable Characteristics

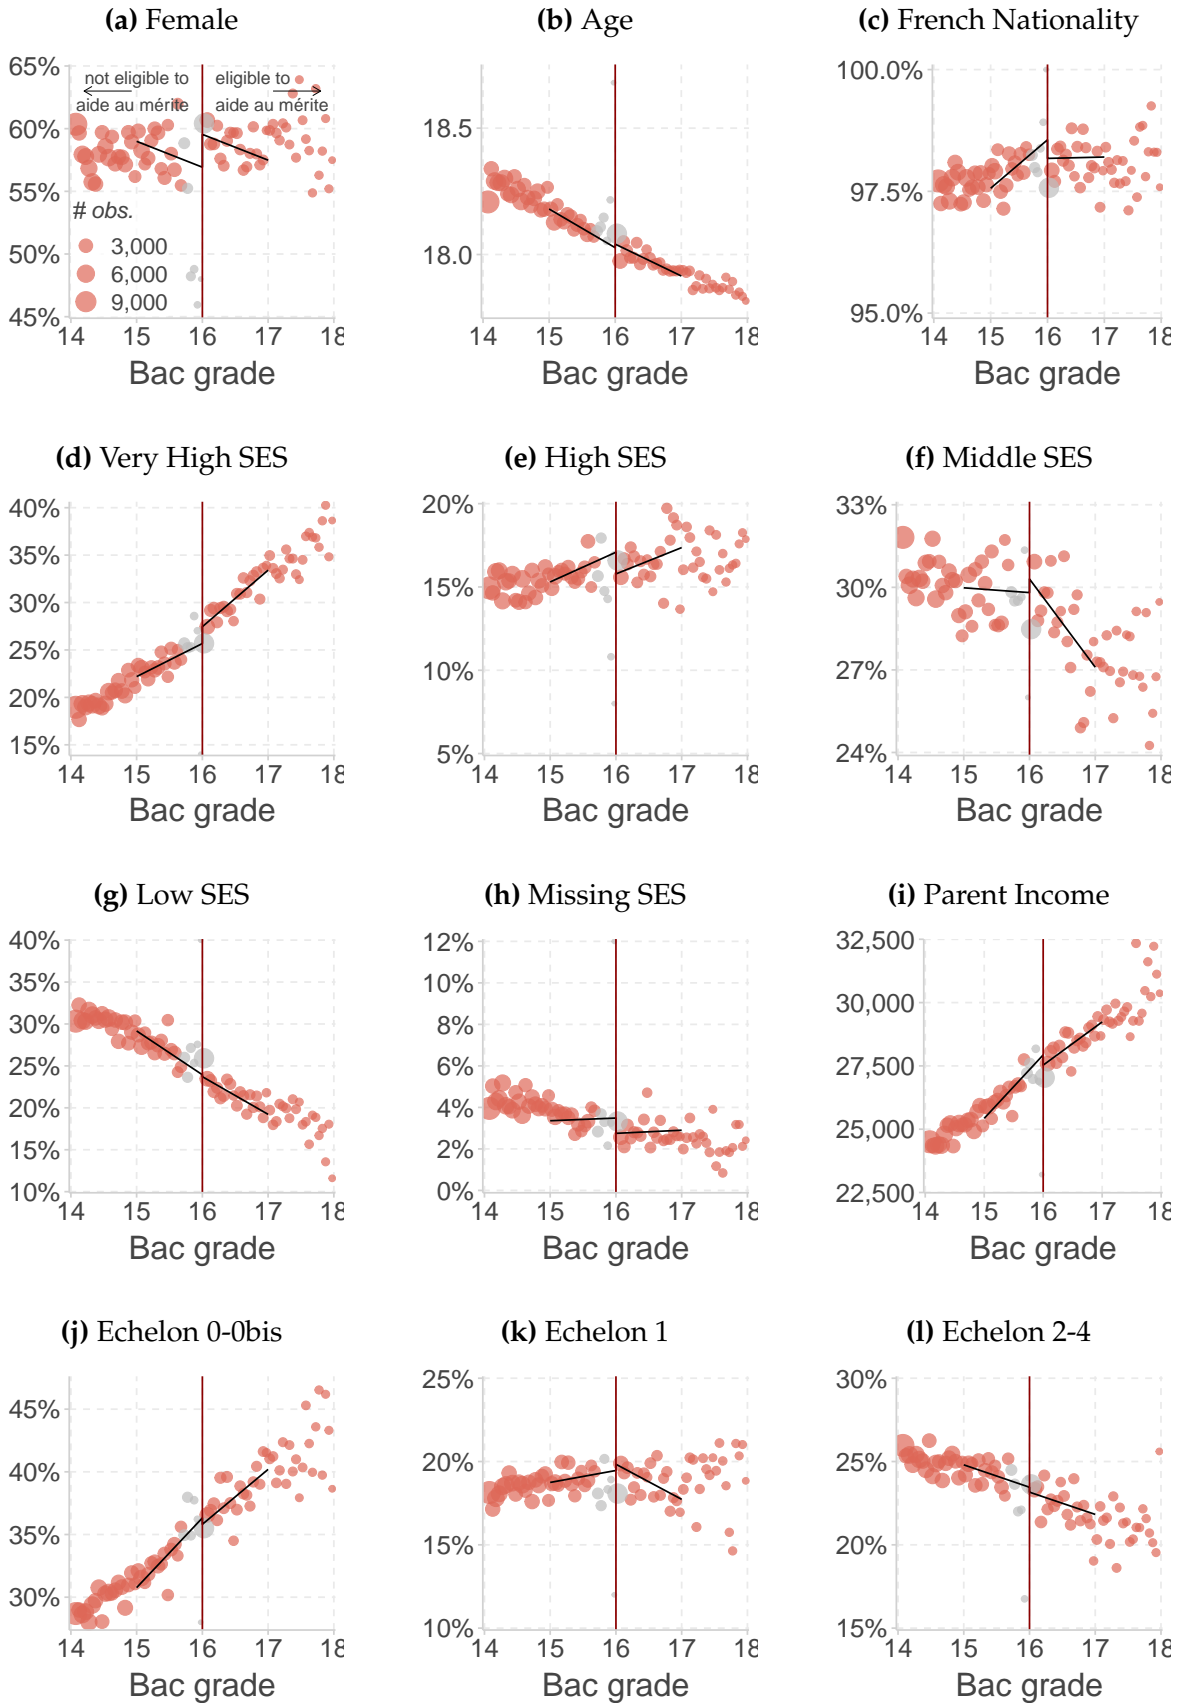
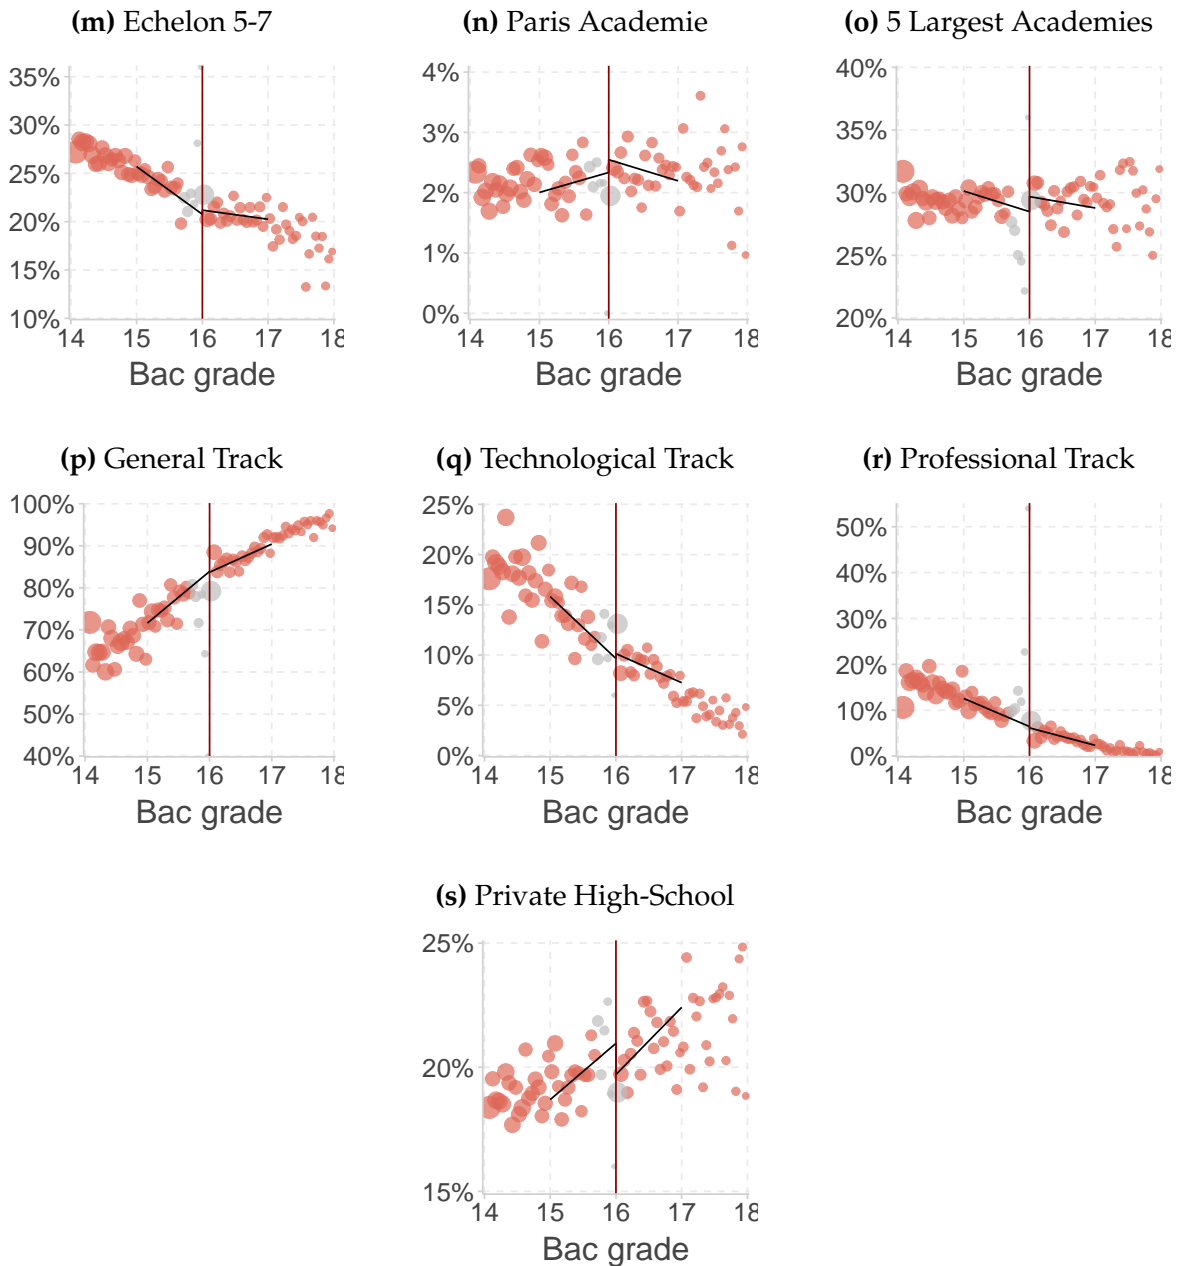



Figure A.6. Discontinuity in Pre-Treatment Observable Characteristics (*continued*)




Notes: This figure shows non-parametric binned scatter plots of the relationship between various student characteristics and Bac grade, around the aide au mérite eligibility threshold (16/20 Bac grade). The student characteristics are reported in the subfigure captions. The sample used corresponds to students from the 2009-2014 Bac cohorts who obtained the Bac, filed a financial aid application and were eligible to a need-based grant. Each red bubble corresponds to the average outcome value for students with Bac grade in $[X, X + 0.05)$, with the size of the bubble corresponding to the number of observations in that grade range. The grey bubbles represent the excluded donut observations, that is observations in $[15.7, 16.05]$. The black fitted lines correspond to local linear regressions with a triangular kernel on each side of the threshold, using the (15, 17) bandwidth, and excluding donut observations. Note that each subfigure has its own y-axis scale.

Figure A.7. Effect of Eligibility to the Aide au Mérite in Bac Year on Degree Choice


(a) Public university (excl. medical degrees)




(b) Medical degrees



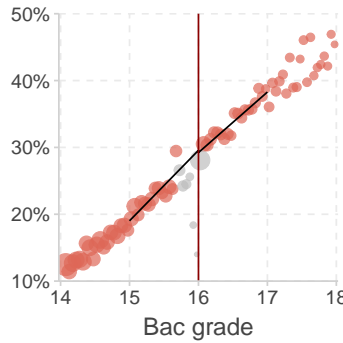
(c) STS




(d) IUT



(e) CPGE




(f) Other



Notes: This figure shows non-parametric binned scatter plots of the relationship between various main enrollment degree types and Bac grade, around the aide au mérite eligibility threshold (16/20 Bac grade). The main enrollment degree types are reported in the subfigure captions. The sample used corresponds to students from the 2009-2014 Bac cohorts who obtained the Bac, filed a financial aid application and were eligible to a need-based grant. Each red bubble corresponds to the average outcome value for students with Bac grade in $[X, X + 0.05]$, with the size of the bubble corresponding to the number of observations in that grade range. The grey bubbles represent the excluded donut observations, that is observations in $[15.7, 16.05]$. The black fitted lines correspond to local linear regressions with a triangular kernel on each side of the threshold, using the (15, 17) bandwidth, and excluding donut observations. Note that each subfigure has its own y-axis scale.


Figure A.8. Effect of Eligibility to the Aide au Mérite on Obtaining a Degree




Notes: This figure shows non-parametric binned scatter plots of the relationship between the probability of obtaining a degree and Bac grade, around the aide au mérite eligibility threshold (16/20 Bac grade). The sample used corresponds to students from the 2009-2014 Bac cohorts who obtained the Bac, filed a financial aid application and were eligible to a need-based grant. Each red bubble corresponds to the average outcome value for students with Bac grade in $[X, X + 0.05)$, with the size of the bubble corresponding to the number of observations in that grade range. The grey bubbles represent the excluded donut observations, that is observations in $[15.7, 16.05]$. The black fitted lines correspond to local linear regressions with a triangular kernel on each side of the threshold, using the (15, 17) bandwidth, and excluding donut observations.

Figure A.9. Effect of Eligibility to the Aide au Mérite on Location of Studies


(a) Paris in Bac Year (Urban Unit)




(b) Paris, Marseille or Lyon in Bac Year (Urban Units)



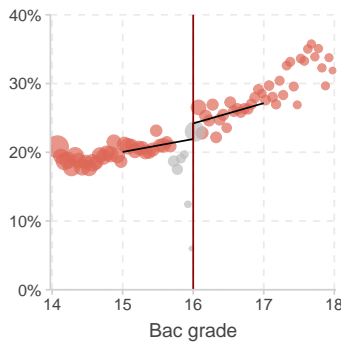
(c) City $\geq 200k$ Inhabitants




(d) Paris in Bac Year + 1 (Urban Unit)



(e) Paris, Marseille or Lyon in Bac Year + 1 (Urban Units)




(f) City $\geq 200k$ Inhabitants in Bac Year + 1



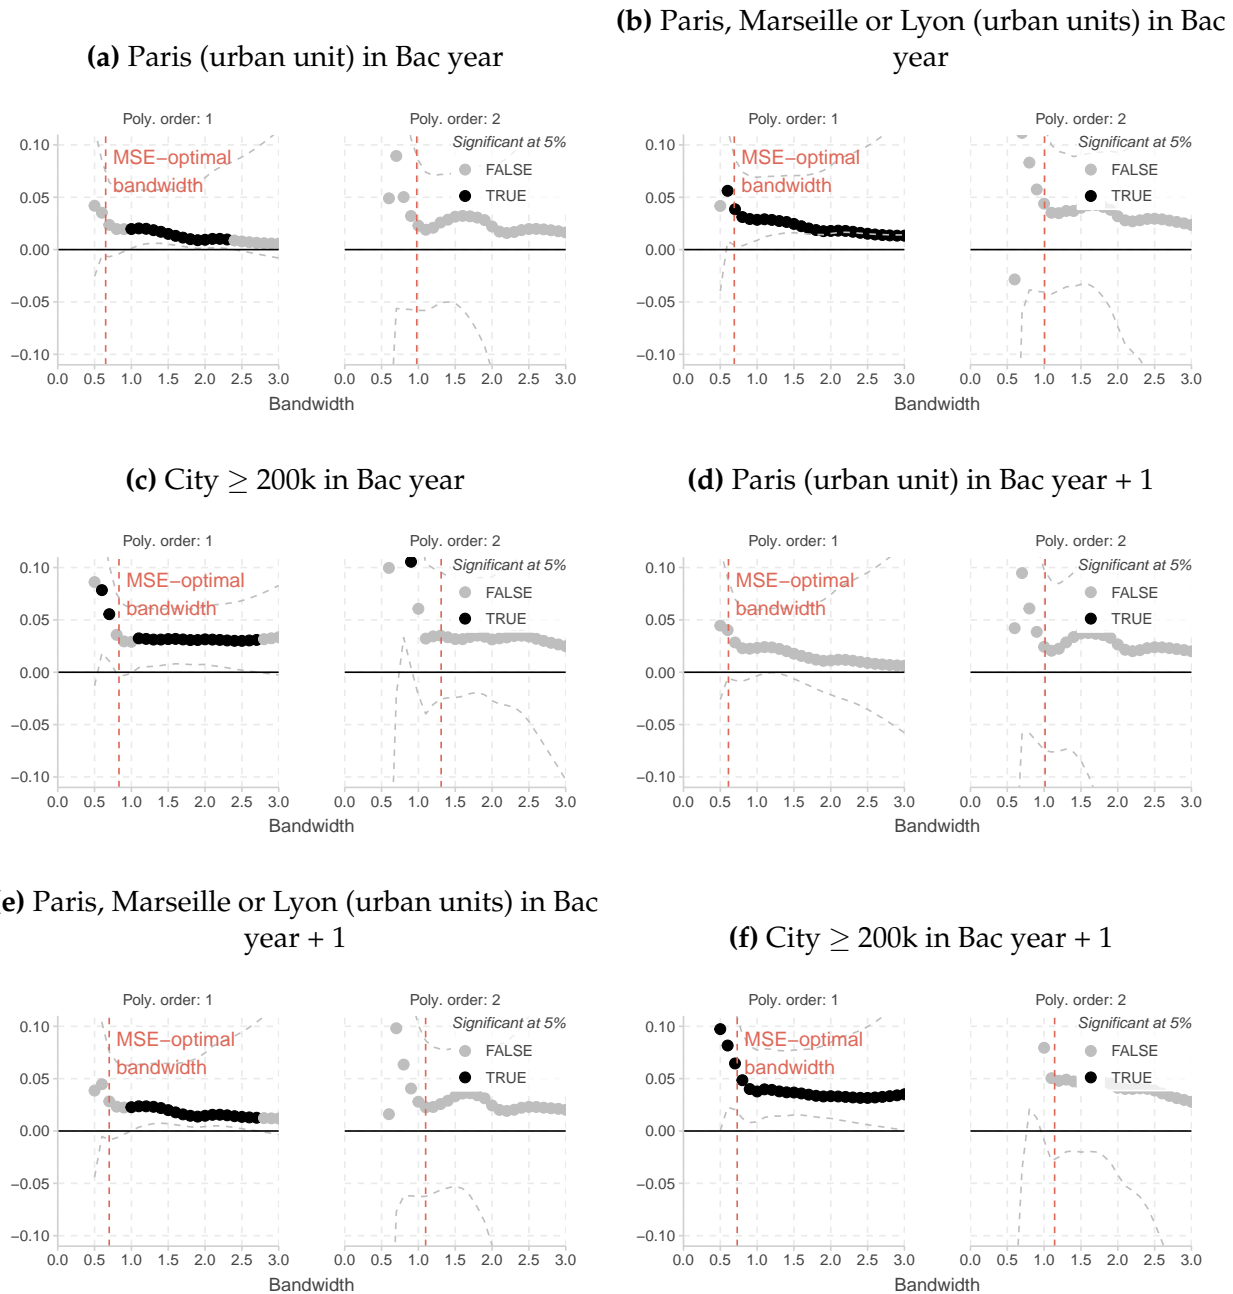
Notes: This figure shows non-parametric binned scatter plots of the relationship between various higher education institution locations and Bac grade, around the aide au mérite eligibility threshold (16/20 Bac grade). The locations are reported in the subfigure captions. The sample used corresponds to students from the 2009-2014 Bac cohorts who obtained the Bac, filed a financial aid application and were eligible to a need-based grant. Each red bubble corresponds to the average outcome value for students with Bac grade in $[X, X + 0.05)$, with the size of the bubble corresponding to the number of observations in that grade range. The grey bubbles represent the excluded donut observations, that is observations in $[15.7, 16.05]$. The black fitted lines correspond to local linear regressions with a triangular kernel on each side of the threshold, using the (15, 17) bandwidth, and excluding donut observations. Note that each subfigure has its own y-axis scale.

Figure A.10. Effect of Eligibility to the Aide au Mérite in Bac Year on Higher Education Outcomes by Bandwidth Size



Notes: This figure shows estimates of the discontinuity in higher education outcomes at the aide au mérite eligibility threshold (16/20 Bac grade), excluding students with Bac grades in [15.7, 16.05], varying the bandwidth over which the estimates are obtained. The higher education outcomes are reported in the subfigure captions. The MSE-optimal bandwidth is denoted with the red dashed line. The dashed grey lines correspond to the robust 95% confidence intervals, where the inference bandwidth (b bandwidth in *rdrobust* terminology) over which these are estimated is fixed over the inference bandwidth of the MSE-optimal point estimate.

Figure A.11. Effect of Eligibility to the Aide au Mérite on Location of Studies by Bandwidth Size



Notes: This figure shows estimates of the discontinuity in location of studies aide au mérite eligibility threshold (16/20 Bac grade), excluding students with Bac grades in $[15.7, 16.05]$, varying the bandwidth over which the estimates are obtained. The higher education institution location are reported in subfigure captions. The cities with over 200,000 inhabitants are: Paris, Marseille, Lyon, Toulouse, Nice, Nantes, Montpellier, Strasbourg, Bordeaux, Lille, and Rennes. The MSE-optimal bandwidth is denoted with the red dashed line. The dashed grey lines correspond to the robust 95% confidence intervals, where the inference bandwidth (b bandwidth in *rdrobust* terminology) over which these are estimated is fixed over the inference bandwidth of the MSE-optimal point estimate.

B. Appendix Tables

Table B.1: Combinations of Parent Income and Disadvantage Points for each Need-Based Grant Echelon for the 2009-10 Academic Year

Echelon # → Points ↓	0	1	2	3	4	5	6
0	32,440	22,060	17,830	15,750	13,710	11,710	7,390
1	36,040	24,510	19,810	17,500	15,230	13,010	8,210
2	39,650	26,960	21,790	19,250	16,760	14,310	9,030
⋮	⋮	⋮	⋮	⋮	⋮	⋮	⋮
16	90,110	61,280	49,530	43,750	38,080	32,530	20,530
17	93,720	63,730	51,510	45,500	39,610	33,830	21,350

Notes: See *Arrêté du 18 août 2009 fixant les plafonds de ressources relatifs aux bourses d'enseignement supérieur du ministère de l'enseignement supérieur et de la recherche pour l'année universitaire 2009-2010.*

Table B.2: Need-Based Grants Annual Amounts by Echelon

Echelon	2009-10	2010-11	2011-12	2012-13	2013-14	2014-15	Aide au mérite (% 2009-10)
0	Exemption from tuition and student social security fees						-
0 bis	-	-	-	-	1,000	1,007	-
1	1,445	1,525	1,606	1,640	1,653	1,665	125
2	2,177	2,298	2,419	2,470	2,490	2,507	83
3	2,790	2,945	3,100	3,165	3,190	3,212	65
4	3,401	3,590	3,779	3,858	3,889	3,916	53
5	3,905	4,122	4,339	4,430	4,465	4,496	46
6	4,140	4,370	4,600	4,697	4,735	4,768	44
7	-	-	-	-	5,500	5,539	-

Notes: Amounts are not adjusted for inflation. All echelons above 0 are also exempt from tuition and student social security fees. Students in the academic regions of Créteil, Paris and Versailles received an additional 153 euros annually. Aide au mérite amount = 1,800 euros per year.

Table B.3: Number of Observations at Each Sample Restriction

Restriction ↓ / Bac cohort →	2009	2010	2011	2012	2013	2014	2009-2014	%
Raw number of obs.	652,109	648,555	690,726	753,742	712,160	745,818	4,203,110	100
+ Obtained the Bac in June session	550,483	544,209	581,087	626,263	608,782	646,121	3,556,945	84.63
+ Unique and non-missing student identifier	525,552	520,859	556,712	588,513	578,586	606,892	3,377,114	94.94
+ Obtained the Bac only once over the period	523,691	519,019	554,809	586,409	576,870	605,564	3,366,362	99.68
+ Bac grade not missing	522,916	518,505	554,404	586,026	576,479	605,384	3,363,714	99.92
+ Eligible to a need-based grant in Bac year	169,660	170,917	183,990	193,853	190,941	192,297	1,101,658	32.75

Notes: This table shows the number of observations for each Bac cohort and for the full sample (2009-2014) at each sample restriction mentioned in Section 3.2.

Table B.4: Descriptive Statistics

	Full sample <i>Need-based grant eligibles in Bac year, 2009-2014 Bac cohorts</i>	Aide au merite eligibles in Bac year <i>Bac grade ≥ 16</i>	RD sample <i>Bac grade $\in (15, 17)$</i>
	(1)	(2)	(3)
<i>Panel A. Socio-demographic characteristics</i>			
Female (%)	56.6	59.3	58.5
Age at Bac (mean)	18.5	18.0	18.1
Parents' taxable income (median; euros)	21,492	28,910	27,133
Very high SES (%)	14.4	31.9	26.2
High or middle SES (%)	43.2	44.6	45.6
Low SES (%)	37.3	20.8	25.0
Missing SES (%)	5.1	2.6	3.2
<i>Panel B. Academic characteristics</i>			
$\geq 16/20$ at Bac (%)	5.0	100.0	47.7
Bac grade (mean)	12.0	16.8	15.8
General high-school track (%)	59.5	88.9	80.4
Technologic high-school track (%)	26.5	7.5	11.6
Professional high-school track (%)	14.0	3.6	8.0
Private high-school (%)	15	21	20
<i>Panel C. Financial aid</i>			
Eligible to aide au merite in Bac year	55,347	55,347	35,879
Eligible to aide au merite in Bac year + 1	49,122	49,122	31,767
Eligible to aide au merite in Bac year + 2	42,754	42,754	27,431
Echelon 0-0 bis (%)	23.1	39.1	35.1
Echelon 1 (%)	16.5	18.9	18.8
Echelon 2-4 (%)	25.9	22.0	23.5
Echelon 5-7 (%)	34.5	20.0	22.6
<i>Panel D. Higher education outcomes</i>			
Enrolled in Bac year (%)	90.4	95.7	94.5
<i>Among enrolled:</i>			
Public university (%)	53.1	39.0	42.1
Vocational degree (STS) (%)	24.6	6.5	11.8
Technical degree (IUT) (%)	12.8	5.7	9.6
Academic preparatory classes (CPGE) (%)	6.8	37.6	28.9
Other institutions (%)	2.7	11.3	7.6
Enrolled in 2nd year in Bac year + 1 (%)	49.4	78.5	73.8
Enrolled in 3rd year in Bac year + 2 (%)	28.2	66.3	58.3
Obtained a degree (2009-2019) (%)	45.9	58.1	60.3
Observations	1,101,658	55,347	75,188

Notes: This table shows descriptive statistics for three samples: (1) the full sample, i.e., students from the 2009-14 Bac cohorts who are eligible to a need-based grant in their Bac year, (2) the aide au mérite eligibles in their Bac year, i.e., students from the full sample who obtained at least 16/20 at the Bac, and (3) the RD sample, i.e., students from the full sample who obtained between 15 and 17 at the Bac. Since 13% of students in the full sample have multiple enrollments in their Bac year, I follow [Bonneau et al. \(2021\)](#)'s ranking across enrollments to assign them a main enrollment. This ranking is based on knowledge of the French higher education system.

Table B.5: Discontinuity in Predicted Outcomes at Aide au Mérite Eligibility Treshold (16/20 Bac Grade)

	No Donut		Donut [15.7, 16.05]	
	MSE-Optimal (1)	Bandwidth: (15, 17) (2)	MSE-Optimal (3)	(15, 17) (4)
Predicted enrollment in Bac Year (adj. R2 = 0.04)	0.004*** [0, 0.01]	0.006*** [0, 0.01]	0.002 [0, 0.01]	0 [0, 0]
Predicted any enrollment (adj. R2 = 0.04)	0.012*** [0.01, 0.01]	0.005*** [0, 0.01]	0.001 [0, 0]	0 [0, 0]
Predicted persistence in 2nd year in Bac year + 1 (adj. R2 = 0.04)	0.011*** [0.01, 0.02]	0.003** [0, 0.01]	-0.004 [-0.01, 0]	-0.002 [-0.01, 0]
Predicted persistence in 2nd Year (adj. R2 = 0.09)	0.056*** [0.04, 0.07]	0.011*** [0.01, 0.02]	-0.001 [-0.01, 0.01]	-0.002 [-0.01, 0]
Predicted persistence in 3rd year in Bac year + 2 (adj. R2 = 0.11)	0.052*** [0.04, 0.07]	0.009*** [0, 0.01]	-0.002 [-0.01, 0.01]	-0.001 [-0.01, 0]
Predicted persistence in 3rd year (adj. R2 = 0.22)	0.105*** [0.09, 0.13]	0.018*** [0.01, 0.03]	-0.001 [-0.02, 0.01]	-0.002 [-0.01, 0.01]
Predicted number of years in HE (adj. R2 = 0.23)	0.487*** [0.39, 0.61]	0.084*** [0.05, 0.12]	0.002 [-0.09, 0.08]	-0.007 [-0.05, 0.04]
Predicted highest level of study attained (adj. R2 = 0.25)	0.364*** [0.29, 0.46]	0.063*** [0.04, 0.09]	-0.005 [-0.07, 0.05]	-0.007 [-0.04, 0.03]
Predicted degree obtention (adj. R2 = 0.16)	0.108*** [0.09, 0.13]	0.019*** [0.01, 0.03]	0 [-0.01, 0.01]	-0.001 [-0.01, 0.01]
Predicted enrollment in masters degree (adj. R2 = 0.2)	0.084*** [0.07, 0.11]	0.014*** [0.01, 0.02]	-0.001 [-0.02, 0.01]	-0.001 [-0.01, 0.01]
Predicted enrollment in selective masters degree (adj. R2 = 0.06)	0.003 [0, 0.01]	-0.001 [0, 0]	-0.003 [-0.01, 0]	-0.002 [0, 0]

Notes: This table reports estimates of the discontinuity in predicted outcomes at the aide au mérite eligibility threshold (16/20 Bac grade). Each predicted outcome is reported in the first column's rows, with the prediction regression adjusted R² in parenthesis. The predictors used for the predictions are: female dummy, age, French nationality dummy, SES (5 categories), parent income, need-based grant echelon (4 cat.), education academie (Paris, 5 largest), high school track (3 categories) and private high school dummy (i.e., students' characteristics in Table 3.1). The prediction model includes no interactions and is estimated by OLS. I report full sample estimates ("No Donut") and estimates obtained when excluding students with Bac grades in [15.7, 16.05] ("Donut [15.7, 16.05]"). Moreover, estimates for two different bandwidths are reported: the MSE-optimal and (15, 17) bandwidths. The MSE-optimal bandwidth, obtained using the *rdrobust* R package, varies across each outcome. For MSE-optimal bandwidth estimates, the ranges in brackets correspond to associated robust 95% confidence intervals, while they correspond to associated conventional 95% confidence intervals for (15, 17) bandwidth estimates. For example, column (1) indicates that, using the MSE-optimal bandwidth and keeping all observations, the estimated discontinuity in the predicted enrollment in Bac year around 16/20 is 0.4 percentage points. This discontinuity estimate is 0.2 percentage points in the donut specification (column (3)). Statistical significance is computed based on the relevant p-value and ***, **, and * indicate significance at 1, 5, and 10%, respectively.

Table B.6: Placebo Analysis using (15, 17) Bandwidth

	Enrollment in Bac Year (1)	Enrollment in 2nd Year in Bac Year + 1 (2)	Enrollment in 3rd Year in Bac Year + 2 (3)	Number of Years in HE (4)	Highest Level of Study Attained (5)	Obtaining a Degree (6)
<i>Panel A. Grade 15</i>						
No donut	0.007** [0, 0.01]	0.004 [-0.01, 0.02]	0.008 [0, 0.02]	0.047* [-0.01, 0.1]	0.04** [0, 0.08]	0.001 [-0.01, 0.01]
Donut [14.7, 15.05]	-0.003 [-0.01, 0.01]	0 [-0.02, 0.02]	0.008 [-0.01, 0.03]	0.045 [-0.04, 0.13]	0.051 [-0.01, 0.11]	0.003 [-0.02, 0.02]
<i>Panel B. Grade 14</i>						
No donut	0.012*** [0.01, 0.02]	0.025*** [0.01, 0.03]	0.051*** [0.04, 0.06]	0.289*** [0.24, 0.34]	0.245*** [0.21, 0.28]	0.056*** [0.05, 0.07]
Donut [13.7, 14.05]	0.003 [-0.01, 0.01]	0.001 [-0.01, 0.02]	0.002 [-0.01, 0.02]	0.049 [-0.02, 0.12]	0.05* [0, 0.1]	0.001 [-0.01, 0.02]
<i>Panel C. Grade 16 for non-eligibles to need-based grants</i>						
No donut	0.027*** [0.02, 0.04]	0.029*** [0.02, 0.04]	0.035*** [0.02, 0.05]	0.183*** [0.11, 0.25]	0.164*** [0.11, 0.21]	0.041*** [0.03, 0.05]
Donut [15.7, 16.05]	-0.013* [-0.03, 0]	-0.005 [-0.02, 0.01]	-0.001 [-0.02, 0.02]	-0.053 [-0.15, 0.05]	-0.057 [-0.13, 0.01]	-0.001 [-0.02, 0.02]

Notes: This table reports estimates of the discontinuity in higher education outcomes, using the (15, 17) bandwidth, for three placebo grade thresholds: grade 15 (panel A), grade 14 (Panel B), and grade 16 for students not eligible to a need-based grant (panel C). Each higher education outcome is reported in the column headers. The grade 15 and grade 14 placebos are estimated over the the sample of students from the 2009-2014 Bac cohorts who obtained the Bac, filed a financial aid application in their Bac year and were eligible to a need-based grant in their Bac year. The grade 16 placebo is estimates on the sample of students from the same Bac cohorts who were not eligible to a need-based grant in their Bac year. In all three cases, students on both sides of placebo grade threshold are not eligible to different amounts of financial aid. I report full sample estimates (“No Donut”) and estimates obtained when excluding students with Bac grades in [placebo – 0.3, placebo + 0.05] (“Donut [...]”). Statistical significance is computed based on the robust p-value and ***, **, and * indicate significance at 1, 5, and 10%, respectively.

Table B.7: Test of Whether aide au mérite Eligibility Had Effects for Students with No or Small Amounts of Cash Allowance as Part of their Need-Based Grant - Parent Income Within 5% of Parent Income Threshold (2009-2012)

	Enrollment in Bac Year (1)	Enrollment in 2nd Year 2 Years After Bac (2)	Enrollment in 3rd Year 3 Years After Bac (3)	Number of Years in HE (4)	Highest Level of Study Attained (5)	Obtaining Degree (6)
<i>Panel A. Echelon 0</i>						
Mean [15.5, 15.7)	0.92	0.7	0.58	5.51	4.47	0.64
MSE-optimal bandwidth	0.055	0.04	0.074	-0.645	-0.266	0.055
(15, 17) bandwidth	[-0.04, 0.16]	[-0.17, 0.2]	[-0.14, 0.26]	[-2.04, 0.68]	[-1.01, 0.46]	[-0.15, 0.26]
	0.056 [-0.02, 0.13]	0.035 [-0.11, 0.18]	0.075 [-0.09, 0.24]	-0.33 [-1.09, 0.42]	-0.224 [-0.68, 0.23]	0.054 [-0.11, 0.22]
<i>Panel B. Echelons 1</i>						
Mean [15.5, 15.7)	0.94	0.76	0.59	5.38	4.37	0.61
MSE-optimal bandwidth	0.004	0.076	0.209**	0.139	0.244	0.065
(15, 17) bandwidth	[-0.07, 0.08]	[-0.14, 0.33]	[0.03, 0.44]	[-0.97, 1.13]	[-0.36, 0.79]	[-0.08, 0.25]
	0.015 [-0.06, 0.09]	0.067 [-0.06, 0.2]	0.145** [0, 0.29]	0.038 [-0.67, 0.75]	0.181 [-0.25, 0.61]	0.078 [-0.07, 0.23]

Notes: This table reports estimates of the discontinuity in higher education outcomes at the aide au mérite eligibility threshold (16/20 Bac grade), excluding students with Bac grades in [15.7, 16.05] for the exact same specification as Table 3.3 except the sample is restricted to students whose parent income is within 5% of the echelon 1 income threshold in 2009-2012. Statistical significance is computed based on the relevant p-value and ***, **, and * indicate significance at 1, 5, and 10%, respectively.

C. Detailed Regression Tables

Table C.1: Effect of Eligibility to the Aide au Mérite in Bac Year on Enrollment in Bac Year

	<i>First order</i>			<i>Second order</i>		
	(1)	(2)	(3)	(4)	(5)	(6)
<i>Panel A: MSE-optimal bandwidth</i>						
Eligibility	0.004	-0.005	-0.003	0.009	-0.006	-0.009
Robust 95% CI	[0, 0.01]	[-0.02, 0.01]	[-0.02, 0.01]	[0, 0.02]	[-0.03, 0.01]	[-0.03, 0.01]
Robust p-value	0.361	0.294	0.539	0.164	0.433	0.261
# obs. left	58,463	33,599	25,535	77,108	56,892	77,447
# obs. right	40,919	24,826	22,272	44,193	30,778	34,343
Bandwidth	(14.73, 17.27)	(15.04, 16.96)	(15.17, 16.83)	(14.48, 17.52)	(14.69, 17.31)	(14.44, 17.56)
<i>Panel B: (15, 17) bandwidth</i>						
Eligibility	0.007	-0.005	-0.005	0.026	0.005	0.002
Conventional 95% CI	[0, 0.02]	[-0.02, 0.01]	[-0.02, 0.01]	[0.01, 0.04]	[-0.03, 0.04]	[-0.03, 0.04]
Conventional p-value	0.136	0.402	0.410	0.007	0.785	0.934
# obs. left	39,274	35,234	35,234	39,274	35,234	35,234
# obs. right	35,879	25,357	25,357	35,879	25,357	25,357
Bandwidth	(15, 17)	(15, 17)	(15, 17)	(15, 17)	(15, 17)	(15, 17)
Poly. order	1	1	1	2	2	2
Donut		✓	✓		✓	✓
Controls			✓			✓
Mean [15.5, 15.7)	0.944	0.944	0.944	0.944	0.944	0.944

Notes: This table reports estimates of the discontinuity in enrollment in Bac year at the aide au mérite eligibility threshold (16/20 Bac grade). Panel A reports estimates using the MSE-optimal bandwidth (obtained from the *rdrobust* R package). Panel B reports estimates using the (15, 17) bandwidth. For MSE-optimal bandwidth estimates, the reported confidence intervals correspond to associated robust 95% confidence intervals, while they correspond to associated conventional 95% confidence intervals for (15, 17) bandwidth estimates. For each panel, the first row reports the discontinuity estimate in the outcome at the 16/20 Bac grade threshold. Columns (1)-(3) report estimates obtained from local linear regressions (triangular kernel) while columns (4)-(6) report estimates from regressions including a quadratic of the running variable. Columns (1) and (4) report estimates on the full sample, columns (2) and (5) estimates excluding students with Bac grades in [15.7, 16.05] ("Donut"), and columns (3) and (6) estimates excluding students with Bac grades in [15.7, 16.05] and including control variables (gender, age, SES, Bac track, and Bac cohort). For example, column (1) indicates that, using the MSE-optimal bandwidth and including all students from the full sample, the estimated discontinuity in enrollment in Bac year around 16/20 is 0.4 percentage points, with the share of such students enrolled among students with Bac grade in [15.5, 15.7) being 94.4%. This discontinuity estimate is 0.7 percentage points when using the (15, 17) bandwidth.

Table C.2: Effect of Eligibility to the Aide au Mérite in Bac Year on Enrollment at Least Once

	<i>First order</i>			<i>Second order</i>		
	(1)	(2)	(3)	(4)	(5)	(6)
<i>Panel A: MSE-optimal bandwidth</i>						
Eligibility	0.002	-0.001	-0.002	0.018	0.011	0.007
Robust 95% CI	[0, 0.01]	[-0.01, 0.01]	[-0.01, 0.01]	[0.01, 0.03]	[-0.01, 0.04]	[-0.01, 0.03]
Robust p-value	0.388	0.581	0.589	0.006	0.299	0.416
# obs. left	60,206	43,505	57,326	44,392	40,247	50,907
# obs. right	41,228	27,783	31,048	37,676	26,825	29,561
Bandwidth	(14.71, 17.29)	(14.89, 17.11)	(14.68, 17.32)	(14.92, 17.08)	(14.93, 17.07)	(14.77, 17.23)
<i>Panel B: (15, 17) bandwidth</i>						
Eligibility	0.006	-0.001	0	0.02	0.015	0.013
Conventional 95% CI	[0, 0.01]	[-0.01, 0.01]	[-0.01, 0.01]	[0.01, 0.03]	[-0.01, 0.04]	[-0.01, 0.04]
Conventional p-value	0.078	0.887	0.936	0.006	0.249	0.313
# obs. left	39,274	35,234	35,234	39,274	35,234	35,234
# obs. right	35,879	25,357	25,357	35,879	25,357	25,357
Bandwidth	(15, 17)	(15, 17)	(15, 17)	(15, 17)	(15, 17)	(15, 17)
Poly. order	1	1	1	2	2	2
Donut		✓	✓		✓	✓
Controls			✓			✓
Mean [15.5, 15.7)	0.973	0.973	0.973	0.973	0.973	0.973

Notes: This table reports estimates of the discontinuity in enrollment at least once at the aide au mérite eligibility threshold (16/20 Bac grade). Panel A reports estimates using the MSE-optimal bandwidth (obtained from the *rdrobust* R package). Panel B reports estimates using the (15, 17) bandwidth. For MSE-optimal bandwidth estimates, the reported confidence intervals correspond to associated robust 95% confidence intervals, while they correspond to associated conventional 95% confidence intervals for (15, 17) bandwidth estimates. For each panel, the first row reports the discontinuity estimate in the outcome at the 16/20 Bac grade threshold. Columns (1)-(3) report estimates obtained from local linear regressions (triangular kernel) while columns (4)-(6) report estimates from regressions including a quadratic of the running variable. Columns (1) and (4) report estimates on the full sample, columns (2) and (5) estimates excluding students with Bac grades in [15.7, 16.05] (“Donut”), and columns (3) and (6) estimates excluding students with Bac grades in [15.7, 16.05] and including control variables (gender, age, SES, Bac track, and Bac cohort). See Table C.1’s notes for an example of how to read the estimates.

Table C.3: Effect of Eligibility to the Aide au Mérite in Bac Year on Enrollment in 2nd Year in Bac Year + 1

	First order			Second order		
	(1)	(2)	(3)	(4)	(5)	(6)
<i>Panel A: MSE-optimal bandwidth</i>						
Eligibility	0.001	-0.001	-0.011	0.052***	0.076	0.073
Robust 95% CI	[-0.02, 0.02]	[-0.04, 0.02]	[-0.03, 0.01]	[0.02, 0.1]	[-0.02, 0.2]	[-0.03, 0.2]
Robust p-value	0.854	0.682	0.198	0.005	0.107	0.136
# obs. left	51,813	30,331	58,868	33,034	27,203	26,042
# obs. right	39,597	23,948	31,161	34,275	23,175	22,620
Bandwidth	(14.82, 17.18)	(15.09, 16.91)	(14.66, 17.34)	(15.1, 16.9)	(15.14, 16.86)	(15.16, 16.84)
<i>Panel B: (15, 17) bandwidth</i>						
Eligibility	0.009	-0.006	-0.005	0.047***	0.055	0.056
Conventional 95% CI	[-0.01, 0.03]	[-0.03, 0.02]	[-0.03, 0.02]	[0.01, 0.08]	[-0.02, 0.13]	[-0.01, 0.13]
Conventional p-value	0.267	0.593	0.668	0.005	0.123	0.114
# obs. left	39,274	35,234	35,234	39,274	35,234	35,234
# obs. right	35,879	25,357	25,357	35,879	25,357	25,357
Bandwidth	(15, 17)	(15, 17)	(15, 17)	(15, 17)	(15, 17)	(15, 17)
Poly. order	1	1	1	2	2	2
Donut		✓	✓		✓	✓
Controls			✓			✓
Mean [15.5, 15.7)	0.726	0.726	0.726	0.726	0.726	0.726

Notes: This table reports estimates of the discontinuity in enrollment in 2nd year in Bac year + 1 at the aide au mérite eligibility threshold (16/20 Bac grade). Panel A reports estimates using the MSE-optimal bandwidth (obtained from the *rdrobust* R package). Panel B reports estimates using the (15, 17) bandwidth. For MSE-optimal bandwidth estimates, the reported confidence intervals correspond to associated robust 95% confidence intervals, while they correspond to associated conventional 95% confidence intervals for (15, 17) bandwidth estimates. For each panel, the first row reports the discontinuity estimate in the outcome at the 16/20 Bac grade threshold. Columns (1)-(3) report estimates obtained from local linear regressions (triangular kernel) while columns (4)-(6) report estimates from regressions including a quadratic of the running variable. Columns (1) and (4) report estimates on the full sample, columns (2) and (5) estimates excluding students with Bac grades in [15.7, 16.05] ("Donut"), and columns (3) and (6) estimates excluding students with Bac grades in [15.7, 16.05] and including control variables (gender, age, SES, Bac track, and Bac cohort). See Table C.1's notes for an example of how to read the estimates. Statistical significance is computed based on the relevant p-value and ***, **, and * indicate significance at 1, 5, and 10%, respectively.

Table C.4: Effect of Eligibility to the Aide au Mérite in Bac Year on Enrollment in 2nd Year at Least Once

	<i>First order</i>			<i>Second order</i>		
	(1)	(2)	(3)	(4)	(5)	(6)
<i>Panel A: MSE-optimal bandwidth</i>						
Eligibility	0.001	-0.018	-0.018*	0.026**	-0.025	-0.025
Robust 95% CI	[-0.01, 0.01]	[-0.05, 0]	[-0.05, 0]	[0.01, 0.05]	[-0.07, 0.03]	[-0.07, 0.02]
Robust p-value	0.787	0.101	0.093	0.010	0.404	0.300
# obs. left	47,870	20,012	20,807	47,545	36,991	40,352
# obs. right	38,342	20,238	20,263	38,305	26,327	27,154
Bandwidth	(14.88, 17.12)	(15.27, 16.73)	(15.25, 16.75)	(14.88, 17.12)	(14.98, 17.02)	(14.92, 17.08)
<i>Panel B: (15, 17) bandwidth</i>						
Eligibility	0.004	-0.016**	-0.015**	0.032***	-0.023	-0.026
Conventional 95% CI	[-0.01, 0.02]	[-0.03, 0]	[-0.03, 0]	[0.01, 0.06]	[-0.07, 0.02]	[-0.07, 0.02]
Conventional p-value	0.459	0.036	0.046	0.007	0.332	0.266
# obs. left	39,274	35,234	35,234	39,274	35,234	35,234
# obs. right	35,879	25,357	25,357	35,879	25,357	25,357
Bandwidth	(15, 17)	(15, 17)	(15, 17)	(15, 17)	(15, 17)	(15, 17)
Poly. order	1	1	1	2	2	2
Donut		✓	✓		✓	✓
Controls			✓			✓
Mean [15.5, 15.7)	0.905	0.905	0.905	0.905	0.905	0.905

Notes: This table reports estimates of the discontinuity in enrollment in 2nd year at least once at the aide au mérite eligibility threshold (16/20 Bac grade). Panel A reports estimates using the MSE-optimal bandwidth (obtained from the *rdrobust* R package). Panel B reports estimates using the (15, 17) bandwidth. For MSE-optimal bandwidth estimates, the reported confidence intervals correspond to associated robust 95% confidence intervals, while they correspond to associated conventional 95% confidence intervals for (15, 17) bandwidth estimates. For each panel, the first row reports the discontinuity estimate in the outcome at the 16/20 Bac grade threshold. Columns (1)-(3) report estimates obtained from local linear regressions (triangular kernel) while columns (4)-(6) report estimates from regressions including a quadratic of the running variable. Columns (1) and (4) report estimates on the full sample, columns (2) and (5) estimates excluding students with Bac grades in [15.7, 16.05] (“Donut”), and columns (3) and (6) estimates excluding students with Bac grades in [15.7, 16.05] and including control variables (gender, age, SES, Bac track, and Bac cohort). See Table C.1’s notes for an example of how to read the estimates. Statistical significance is computed based on the relevant p-value and ***, **, and * indicate significance at 1, 5, and 10%, respectively.

Table C.5: Effect of Eligibility to the Aide au Mérite in Bac Year on Enrollment in 3rd Year in Bac Year + 2

	<i>First order</i>			<i>Second order</i>		
	(1)	(2)	(3)	(4)	(5)	(6)
<i>Panel A: MSE-optimal bandwidth</i>						
Eligibility	0.021**	0.015	0.01	0.101***	0.088*	0.066
Robust 95% CI	[0.01, 0.04]	[-0.02, 0.05]	[-0.02, 0.04]	[0.07, 0.15]	[-0.01, 0.21]	[-0.02, 0.17]
Robust p-value	0.011	0.442	0.566	0.000	0.077	0.110
# obs. left	54,947	27,203	30,924	39,216	28,994	33,599
# obs. right	40,083	23,175	24,289	35,877	23,753	24,826
Bandwidth	(14.78, 17.22)	(15.13, 16.87)	(15.08, 16.92)	(15.01, 16.99)	(15.11, 16.89)	(15.04, 16.96)
<i>Panel B: (15, 17) bandwidth</i>						
Eligibility	0.033***	0.006	0.006	0.1***	0.069*	0.061
Conventional 95% CI	[0.01, 0.05]	[-0.02, 0.03]	[-0.02, 0.03]	[0.06, 0.14]	[-0.01, 0.15]	[-0.01, 0.14]
Conventional p-value	0.000	0.639	0.639	0.000	0.084	0.115
# obs. left	39,274	35,234	35,234	39,274	35,234	35,234
# obs. right	35,879	25,357	25,357	35,879	25,357	25,357
Bandwidth	(15, 17)	(15, 17)	(15, 17)	(15, 17)	(15, 17)	(15, 17)
Poly. order	1	1	1	2	2	2
Donut		✓	✓		✓	✓
Controls			✓			✓
Mean [15.5, 15.7)	0.563	0.563	0.563	0.563	0.563	0.563

Notes: This table reports estimates of the discontinuity in enrollment in 3rd year in Bac year + 2 at the aide au mérite eligibility threshold (16/20 Bac grade). Panel A reports estimates using the MSE-optimal bandwidth (obtained from the *rdrobust* R package). Panel B reports estimates using the (15, 17) bandwidth. For MSE-optimal bandwidth estimates, the reported confidence intervals correspond to associated robust 95% confidence intervals, while they correspond to associated conventional 95% confidence intervals for (15, 17) bandwidth estimates. For each panel, the first row reports the discontinuity estimate in the outcome at the 16/20 Bac grade threshold. Columns (1)-(3) report estimates obtained from local linear regressions (triangular kernel) while columns (4)-(6) report estimates from regressions including a quadratic of the running variable. Columns (1) and (4) report estimates on the full sample, columns (2) and (5) estimates excluding students with Bac grades in [15.7, 16.05] (“Donut”), and columns (3) and (6) estimates excluding students with Bac grades in [15.7, 16.05] and including control variables (gender, age, SES, Bac track, and Bac cohort). See Table C.1’s notes for an example of how to read the estimates. Statistical significance is computed based on the relevant p-value and ***, **, and * indicate significance at 1, 5, and 10%, respectively.

Table C.6: Effect of Eligibility to the Aide au Mérite in Bac Year on Enrollment in 3rd Year at Least Once

	<i>First order</i>			<i>Second order</i>		
	(1)	(2)	(3)	(4)	(5)	(6)
<i>Panel A: MSE-optimal bandwidth</i>						
Eligibility	0.02***	-0.013	-0.012	0.04***	-0.019	-0.012
Robust 95% CI	[0.01, 0.04]	[-0.03, 0.01]	[-0.03, 0.01]	[0.02, 0.07]	[-0.07, 0.02]	[-0.04, 0.01]
Robust p-value	0.003	0.144	0.220	0.000	0.290	0.315
# obs. left	42,906	45,498	42,043	72,766	52,358	79,912
# obs. right	37,331	28,281	27,307	43,499	29,894	34,670
Bandwidth	(14.94, 17.06)	(14.86, 17.14)	(14.9, 17.1)	(14.55, 17.45)	(14.75, 17.25)	(14.41, 17.59)
<i>Panel B: (15, 17) bandwidth</i>						
Eligibility	0.022***	-0.014	-0.014	0.086***	-0.007	-0.016
Conventional 95% CI	[0.01, 0.04]	[-0.03, 0.01]	[-0.03, 0]	[0.06, 0.12]	[-0.07, 0.05]	[-0.07, 0.04]
Conventional p-value	0.002	0.146	0.129	0.000	0.811	0.560
# obs. left	39,274	35,234	35,234	39,274	35,234	35,234
# obs. right	35,879	25,357	25,357	35,879	25,357	25,357
Bandwidth	(15, 17)	(15, 17)	(15, 17)	(15, 17)	(15, 17)	(15, 17)
Poly. order	1	1	1	2	2	2
Donut		✓	✓		✓	✓
Controls			✓			✓
Mean [15.5, 15.7)	0.816	0.816	0.816	0.816	0.816	0.816

Notes: This table reports estimates of the discontinuity in enrollment in 3rd year at least once at the aide au mérite eligibility threshold (16/20 Bac grade). Panel A reports estimates using the MSE-optimal bandwidth (obtained from the *rdrobust* R package). Panel B reports estimates using the (15, 17) bandwidth. For MSE-optimal bandwidth estimates, the reported confidence intervals correspond to associated robust 95% confidence intervals, while they correspond to associated conventional 95% confidence intervals for (15, 17) bandwidth estimates. For each panel, the first row reports the discontinuity estimate in the outcome at the 16/20 Bac grade threshold. Columns (1)-(3) report estimates obtained from local linear regressions (triangular kernel) while columns (4)-(6) report estimates from regressions including a quadratic of the running variable. Columns (1) and (4) report estimates on the full sample, columns (2) and (5) estimates excluding students with Bac grades in [15.7, 16.05] ("Donut"), and columns (3) and (6) estimates excluding students with Bac grades in [15.7, 16.05] and including control variables (gender, age, SES, Bac track, and Bac cohort). See Table C.1's notes for an example of how to read the estimates. Statistical significance is computed based on the relevant p-value and ***, **, and * indicate significance at 1, 5, and 10%, respectively.

Table C.7: Effect of Eligibility to the Aide au Mérite in Bac Year on the Number of Years Enrolled in Higher Education

	<i>First order</i>			<i>Second order</i>		
	(1)	(2)	(3)	(4)	(5)	(6)
<i>Panel A: MSE-optimal bandwidth</i>						
Eligibility	0.084**	-0.007	0.008	0.241***	0.022	-0.07
Robust 95% CI	[0.01, 0.17]	[-0.21, 0.15]	[-0.09, 0.11]	[0.13, 0.4]	[-0.37, 0.43]	[-0.33, 0.15]
Robust p-value	0.022	0.723	0.890	0.000	0.886	0.448
# obs. left	62,376	24,015	53,272	63,390	36,851	51,424
# obs. right	41,624	21,613	30,001	41,921	25,853	29,808
Bandwidth	(14.67, 17.33)	(15.19, 16.81)	(14.75, 17.25)	(14.65, 17.35)	(14.99, 17.01)	(14.76, 17.24)
<i>Panel B: (15, 17) bandwidth</i>						
Eligibility	0.133***	-0.028	-0.019	0.433***	0.037	0.001
Conventional 95% CI	[0.05, 0.22]	[-0.14, 0.09]	[-0.12, 0.08]	[0.27, 0.6]	[-0.31, 0.39]	[-0.31, 0.31]
Conventional p-value	0.002	0.632	0.723	0.000	0.835	0.995
# obs. left	39,274	35,234	35,234	39,274	35,234	35,234
# obs. right	35,879	25,357	25,357	35,879	25,357	25,357
Bandwidth	(15, 17)	(15, 17)	(15, 17)	(15, 17)	(15, 17)	(15, 17)
Poly. order	1	1	1	2	2	2
Donut		✓	✓		✓	✓
Controls			✓			✓
Mean [15.5, 15.7)	5.23	5.23	5.23	5.23	5.23	5.23

Notes: This table reports estimates of the discontinuity in the number of years enrolled in higher education at the aide au mérite eligibility threshold (16/20 Bac grade). Panel A reports estimates using the MSE-optimal bandwidth (obtained from the *rdrobust* R package). Panel B reports estimates using the (15, 17) bandwidth. For MSE-optimal bandwidth estimates, the reported confidence intervals correspond to associated robust 95% confidence intervals, while they correspond to associated conventional 95% confidence intervals for (15, 17) bandwidth estimates. For each panel, the first row reports the discontinuity estimate in the outcome at the 16/20 Bac grade threshold. Columns (1)-(3) report estimates obtained from local linear regressions (triangular kernel) while columns (4)-(6) report estimates from regressions including a quadratic of the running variable. Columns (1) and (4) report estimates on the full sample, columns (2) and (5) estimates excluding students with Bac grades in [15.7, 16.05] (“Donut”), and columns (3) and (6) estimates excluding students with Bac grades in [15.7, 16.05] and including control variables (gender, age, SES, Bac track, and Bac cohort). See Table C.1’s notes for an example of how to read the estimates. Statistical significance is computed based on the relevant p-value and ***, **, and * indicate significance at 1, 5, and 10%, respectively.

Table C.8: Effect of Eligibility to the Aide au Mérite in Bac Year on the Highest Level of Study Attained (in Years)

	<i>First order</i>			<i>Second order</i>		
	(1)	(2)	(3)	(4)	(5)	(6)
<i>Panel A: MSE-optimal bandwidth</i>						
Eligibility	0.168***	-0.038	-0.05	0.324***	0.023	-0.028
Robust 95% CI	[0.1, 0.26]	[-0.18, 0.07]	[-0.15, 0.03]	[0.23, 0.46]	[-0.33, 0.39]	[-0.28, 0.22]
Robust p-value	0.000	0.402	0.165	0.000	0.860	0.806
# obs. left	26,375	21,646	28,994	41,031	28,440	35,176
# obs. right	31,511	20,890	23,753	36,849	23,236	25,355
Bandwidth	(15.22, 16.78)	(15.23, 16.77)	(15.1, 16.9)	(14.98, 17.02)	(15.12, 16.88)	(15.02, 16.98)
<i>Panel B: (15, 17) bandwidth</i>						
Eligibility	0.087***	-0.058	-0.053	0.339***	0.002	-0.032
Conventional 95% CI	[0.03, 0.14]	[-0.13, 0.02]	[-0.12, 0.01]	[0.23, 0.45]	[-0.23, 0.24]	[-0.24, 0.17]
Conventional p-value	0.002	0.136	0.124	0.000	0.987	0.764
# obs. left	39,274	35,234	35,234	39,274	35,234	35,234
# obs. right	35,879	25,357	25,357	35,879	25,357	25,357
Bandwidth	(15, 17)	(15, 17)	(15, 17)	(15, 17)	(15, 17)	(15, 17)
Poly. order	1	1	1	2	2	2
Donut		✓			✓	✓
Controls			✓			✓
Mean [15.5, 15.7]	4.287	4.287	4.287	4.287	4.287	4.287

Notes: This table reports estimates of the discontinuity in the highest level of study attained at the aide au mérite eligibility threshold (16/20 Bac grade). Panel A reports estimates using the MSE-optimal bandwidth (obtained from the *rdrobust* R package). Panel B reports estimates using the (15, 17) bandwidth. For MSE-optimal bandwidth estimates, the reported confidence intervals correspond to associated robust 95% confidence intervals, while they correspond to associated conventional 95% confidence intervals for (15, 17) bandwidth estimates. For each panel, the first row reports the discontinuity estimate in the outcome at the 16/20 Bac grade threshold. Columns (1)-(3) report estimates obtained from local linear regressions (triangular kernel) while columns (4)-(6) report estimates from regressions including a quadratic of the running variable. Columns (1) and (4) report estimates on the full sample, columns (2) and (5) estimates excluding students with Bac grades in [15.7, 16.05] ("Donut"), and columns (3) and (6) estimates excluding students with Bac grades in [15.7, 16.05] and including control variables (gender, age, SES, Bac track, and Bac cohort). See Table C.1's notes for an example of how to read the estimates. Statistical significance is computed based on the relevant p-value and ***, **, and * indicate significance at 1, 5, and 10%, respectively.

Table C.9: Effect of Eligibility to aide au mérite on Degree Quality in Bac Year

	<i>First order</i>			<i>Second order</i>		
	(1)	(2)	(3)	(4)	(5)	(6)
<i>Panel A: MSE-optimal bandwidth</i>						
Eligibility	0.068	-0.012	0.023	0.221***	-0.031	-0.003
Robust 95% CI	[-0.02, 0.18]	[-0.15, 0.08]	[-0.09, 0.1]	[0.09, 0.4]	[-0.18, 0.1]	[-0.12, 0.11]
Robust p-value	0.101	0.531	0.894	0.002	0.571	0.964
# obs. left	21,069	26,671	32,950	26,629	75,120	87,033
# obs. right	28,357	22,174	24,210	31,025	33,352	34,743
Bandwidth	(15.29, 16.71)	(15.12, 16.88)	(15.02, 16.98)	(15.18, 16.82)	(14.4, 17.6)	(14.25, 17.75)
<i>Panel B: (15, 17) bandwidth</i>						
Eligibility	0.02	0.001	0.023	0.126**	-0.138	-0.091
Conventional 95% CI	[-0.04, 0.08]	[-0.09, 0.09]	[-0.06, 0.1]	[0.01, 0.24]	[-0.4, 0.12]	[-0.34, 0.15]
Conventional p-value	0.520	0.976	0.582	0.033	0.296	0.465
# obs. left	36,784	33,000	33,000	36,784	33,000	33,000
# obs. right	34,198	24,211	24,211	34,198	24,211	24,211
Bandwidth	(15, 17)	(15, 17)	(15, 17)	(15, 17)	(15, 17)	(15, 17)
Poly. order	1	1	1	2	2	2
Donut		✓	✓		✓	✓
Controls			✓			✓
Mean [15.5, 15.7)	13.282	13.282	13.282	13.282	13.282	13.282

Notes: This table reports estimates of the discontinuity in degree quality in Bac year at the aide au mérite eligibility threshold (16/20 Bac grade). Degree quality is measured as the median Bac grade among contemporaneously enrolled students. Panel A reports estimates using the MSE-optimal bandwidth (obtained from the *rdrobust* R package). Panel B reports estimates using the (15, 17) bandwidth. For MSE-optimal bandwidth estimates, the reported confidence intervals correspond to associated robust 95% confidence intervals, while they correspond to associated conventional 95% confidence intervals for (15, 17) bandwidth estimates. For each panel, the first row reports the discontinuity estimate in the outcome at the 16/20 Bac grade threshold. Columns (1)-(3) report estimates obtained from local linear regressions (triangular kernel) while columns (4)-(6) report estimates from regressions including a quadratic of the running variable. Columns (1) and (4) report estimates on the full sample, columns (2) and (5) estimates excluding students with Bac grades in [15.7, 16.05] ("Donut"), and columns (3) and (6) estimates excluding students with Bac grades in [15.7, 16.05] and including control variables (gender, age, SES, Bac track, and Bac cohort). See Table C.1's notes for an example of how to read the estimates. Statistical significance is computed based on the relevant p-value and ***, **, and * indicate significance at 1, 5, and 10%, respectively.

Table C.10: Effect of Eligibility to aide au mérite on Degree Quality in Bac Year + 1

	First order			Second order		
	(1)	(2)	(3)	(4)	(5)	(6)
<i>Panel A: MSE-optimal bandwidth</i>						
Eligibility	0.03	0.011	0.022	0.067	-0.017	0.005
Robust 95% CI	[-0.06, 0.14]	[-0.17, 0.14]	[-0.15, 0.16]	[-0.05, 0.22]	[-0.22, 0.18]	[-0.18, 0.18]
Robust p-value	0.480	0.850	0.960	0.206	0.841	0.979
# obs. left	20,397	17,691	17,691	36,962	48,773	50,307
# obs. right	27,723	18,574	18,574	33,953	28,558	28,849
Bandwidth	(15.29, 16.71)	(15.28, 16.72)	(15.28, 16.72)	(15, 17)	(14.73, 17.27)	(14.71, 17.29)
<i>Panel B: (15, 17) bandwidth</i>						
Eligibility	-0.001	0.021	0.034	0.067	-0.042	-0.015
Conventional 95% CI	[-0.06, 0.06]	[-0.07, 0.11]	[-0.05, 0.12]	[-0.05, 0.19]	[-0.3, 0.22]	[-0.26, 0.23]
Conventional p-value	0.964	0.642	0.417	0.266	0.750	0.903
# obs. left	35,505	31,843	31,843	35,505	31,843	31,843
# obs. right	33,483	23,778	23,778	33,483	23,778	23,778
Bandwidth	(15, 17)	(15, 17)	(15, 17)	(15, 17)	(15, 17)	(15, 17)
Poly. order	1	1	1	2	2	2
Donut		✓	✓		✓	✓
Controls			✓			✓
Mean [15.5, 15.7)	13.331	13.331	13.331	13.331	13.331	13.331

Notes: This table reports estimates of the discontinuity in degree quality in Bac year + 1 at the aide au mérite eligibility threshold (16/20 Bac grade). Degree quality is measured as the median Bac grade among contemporaneously enrolled students. Panel A reports estimates using the MSE-optimal bandwidth (obtained from the *rdrobust* R package). Panel B reports estimates using the (15, 17) bandwidth. For MSE-optimal bandwidth estimates, the reported confidence intervals correspond to associated robust 95% confidence intervals, while they correspond to associated conventional 95% confidence intervals for (15, 17) bandwidth estimates. For each panel, the first row reports the discontinuity estimate in the outcome at the 16/20 Bac grade threshold. Columns (1)-(3) report estimates obtained from local linear regressions (triangular kernel) while columns (4)-(6) report estimates from regressions including a quadratic of the running variable. Columns (1) and (4) report estimates on the full sample, columns (2) and (5) estimates excluding students with Bac grades in [15.7, 16.05] ("Donut"), and columns (3) and (6) estimates excluding students with Bac grades in [15.7, 16.05] and including control variables (gender, age, SES, Bac track, and Bac cohort). See Table C.1's notes for an example of how to read the estimates. Statistical significance is computed based on the relevant p-value and ***, **, and * indicate significance at 1, 5, and 10%, respectively.

Table C.11: Effect of Eligibility to aide au mérite on Degree Quality in Bac Year + 2

	<i>First order</i>			<i>Second order</i>		
	(1)	(2)	(3)	(4)	(5)	(6)
<i>Panel A: MSE-optimal bandwidth</i>						
Eligibility	-0.088	-0.018	0	-0.085	-0.009	0.016
Robust 95% CI	[-0.23, 0.03]	[-0.21, 0.15]	[-0.18, 0.17]	[-0.21, 0.04]	[-0.18, 0.14]	[-0.17, 0.19]
Robust p-value	0.132	0.726	0.951	0.186	0.797	0.912
# obs. left	11,901	14,062	14,062	37,078	51,188	44,308
# obs. right	22,997	16,772	16,772	33,746	29,143	27,428
Bandwidth	(15.45, 16.55)	(15.32, 16.68)	(15.31, 16.69)	(14.89, 17.11)	(14.58, 17.42)	(14.71, 17.29)
<i>Panel B: (15, 17) bandwidth</i>						
Eligibility	-0.05	0.019	0.039	-0.094	-0.056	-0.002
Conventional 95% CI	[-0.11, 0.01]	[-0.07, 0.1]	[-0.04, 0.12]	[-0.22, 0.03]	[-0.31, 0.2]	[-0.24, 0.24]
Conventional p-value	0.106	0.667	0.345	0.125	0.667	0.986
# obs. left	31,585	28,287	28,287	31,585	28,287	28,287
# obs. right	31,573	22,518	22,518	31,573	22,518	22,518
Bandwidth	(15, 17)	(15, 17)	(15, 17)	(15, 17)	(15, 17)	(15, 17)
Poly. order	1	1	1	2	2	2
Donut		✓	✓		✓	✓
Controls			✓			✓
Mean [15.5, 15.7]	13.389	13.389	13.389	13.389	13.389	13.389

Notes: This table reports estimates of the discontinuity in degree quality in Bac year + 2 at the aide au mérite eligibility threshold (16/20 Bac grade). Degree quality is measured as the median Bac grade among contemporaneously enrolled students. Panel A reports estimates using the MSE-optimal bandwidth (obtained from the *rdrobust* R package). Panel B reports estimates using the (15, 17) bandwidth. For MSE-optimal bandwidth estimates, the reported confidence intervals correspond to associated robust 95% confidence intervals, while they correspond to associated conventional 95% confidence intervals for (15, 17) bandwidth estimates. For each panel, the first row reports the discontinuity estimate in the outcome at the 16/20 Bac grade threshold. Columns (1)-(3) report estimates obtained from local linear regressions (triangular kernel) while columns (4)-(6) report estimates from regressions including a quadratic of the running variable. Columns (1) and (4) report estimates on the full sample, columns (2) and (5) estimates excluding students with Bac grades in [15.7, 16.05] ("Donut"), and columns (3) and (6) estimates excluding students with Bac grades in [15.7, 16.05] and including control variables (gender, age, SES, Bac track, and Bac cohort). See Table C.1's notes for an example of how to read the estimates. Statistical significance is computed based on the relevant p-value and ***, **, and * indicate significance at 1, 5, and 10%, respectively.

Table C.12: Effect of Eligibility to the Aide au Mérite in Bac Year on Obtaining a Degree

	<i>First order</i>			<i>Second order</i>		
	(1)	(2)	(3)	(4)	(5)	(6)
<i>Panel A: MSE-optimal bandwidth</i>						
Eligibility	0.029***	-0.022	-0.027*	0.089***	-0.039	-0.044*
Robust 95% CI	[0.01, 0.05]	[-0.05, 0]	[-0.06, 0]	[0.06, 0.13]	[-0.1, 0.01]	[-0.11, 0]
Robust p-value	0.001	0.115	0.075	0.000	0.119	0.059
# obs. left	34,282	47,773	36,894	41,031	51,424	51,424
# obs. right	34,315	29,075	25,853	36,849	29,808	29,808
Bandwidth	(15.09, 16.91)	(14.81, 17.19)	(14.99, 17.01)	(14.98, 17.02)	(14.76, 17.24)	(14.77, 17.23)
<i>Panel B: (15, 17) bandwidth</i>						
Eligibility	0.022**	-0.026**	-0.027**	0.094***	-0.011	-0.024
Conventional 95% CI	[0, 0.04]	[-0.05, 0]	[-0.05, 0]	[0.06, 0.13]	[-0.09, 0.06]	[-0.1, 0.05]
Conventional p-value	0.018	0.040	0.028	0.000	0.770	0.527
# obs. left	39,274	35,234	35,234	39,274	35,234	35,234
# obs. right	35,879	25,357	25,357	35,879	25,357	25,357
Bandwidth	(15, 17)	(15, 17)	(15, 17)	(15, 17)	(15, 17)	(15, 17)
Poly. order	1	1	1	2	2	2
Donut		✓	✓		✓	✓
Controls			✓			✓
Mean [15.5, 15.7)	0.622	0.622	0.622	0.622	0.622	0.622

Notes: This table reports estimates of the discontinuity in the probability of obtaining a degree at the aide au mérite eligibility threshold (16/20 Bac grade). Panel A reports estimates using the MSE-optimal bandwidth (obtained from the *rdrobust* R package). Panel B reports estimates using the (15, 17) bandwidth. For MSE-optimal bandwidth estimates, the reported confidence intervals correspond to associated robust 95% confidence intervals, while they correspond to associated conventional 95% confidence intervals for (15, 17) bandwidth estimates. For each panel, the first row reports the discontinuity estimate in the outcome at the 16/20 Bac grade threshold. Columns (1)-(3) report estimates obtained from local linear regressions (triangular kernel) while columns (4)-(6) report estimates from regressions including a quadratic of the running variable. Columns (1) and (4) report estimates on the full sample, columns (2) and (5) estimates excluding students with Bac grades in [15.7, 16.05] (“Donut”), and columns (3) and (6) estimates excluding students with Bac grades in [15.7, 16.05] and including control variables (gender, age, SES, Bac track, and Bac cohort). See Table C.1’s notes for an example of how to read the estimates. Statistical significance is computed based on the relevant p-value and ***, **, and * indicate significance at 1, 5, and 10%, respectively.

Table C.13: Effect of Eligibility to the Aide au Mérite in Bac Year on Enrolling in a Masters Degree

	<i>First order</i>			<i>Second order</i>		
	(1)	(2)	(3)	(4)	(5)	(6)
<i>Panel A: MSE-optimal bandwidth</i>						
Eligibility	0.024***	-0.019	-0.017*	0.098***	-0.014	-0.023
Robust 95% CI	[0.01, 0.04]	[-0.06, 0.01]	[-0.04, 0]	[0.07, 0.14]	[-0.12, 0.09]	[-0.11, 0.05]
Robust p-value	0.003	0.191	0.088	0.000	0.790	0.518
# obs. left	47,545	22,335	43,830	46,416	30,242	33,450
# obs. right	38,305	20,989	27,820	38,233	23,793	24,824
Bandwidth	(14.88, 17.12)	(15.22, 16.78)	(14.88, 17.12)	(14.89, 17.11)	(15.1, 16.9)	(15.04, 16.96)
<i>Panel B: (15, 17) bandwidth</i>						
Eligibility	0.032***	-0.02*	-0.018*	0.117***	-0.018	-0.023
Conventional 95% CI	[0.02, 0.05]	[-0.04, 0]	[-0.04, 0]	[0.08, 0.15]	[-0.09, 0.05]	[-0.09, 0.04]
Conventional p-value	0.000	0.091	0.097	0.000	0.628	0.480
# obs. left	39,274	35,234	35,234	39,274	35,234	35,234
# obs. right	35,879	25,357	25,357	35,879	25,357	25,357
Bandwidth	(15, 17)	(15, 17)	(15, 17)	(15, 17)	(15, 17)	(15, 17)
Poly. order	1	1	1	2	2	2
Donut		✓	✓		✓	✓
Controls			✓			✓
Mean [15.5, 15.7)	0.69	0.69	0.69	0.69	0.69	0.69

Notes: This table reports estimates of the discontinuity in the probability of enrolling in a masters degree at the aide au mérite eligibility threshold (16/20 Bac grade). Masters degrees are defined as degrees for which the final year of study is 4 or 5. Panel A reports estimates using the MSE-optimal bandwidth (obtained from the *rdrobust* R package). Panel B reports estimates using the (15, 17) bandwidth. For MSE-optimal bandwidth estimates, the reported confidence intervals correspond to associated robust 95% confidence intervals, while they correspond to associated conventional 95% confidence intervals for (15, 17) bandwidth estimates. For each panel, the first row reports the discontinuity estimate in the outcome at the 16/20 Bac grade threshold. Columns (1)-(3) report estimates obtained from local linear regressions (triangular kernel) while columns (4)-(6) report estimates from regressions including a quadratic of the running variable. Columns (1) and (4) report estimates on the full sample, columns (2) and (5) estimates excluding students with Bac grades in [15.7, 16.05] ("Donut"), and columns (3) and (6) estimates excluding students with Bac grades in [15.7, 16.05] and including control variables (gender, age, SES, Bac track, and Bac cohort). See Table C.1's notes for an example of how to read the estimates. Statistical significance is computed based on the relevant p-value and ***, **, and * indicate significance at 1, 5, and 10%, respectively.

Table C.14: Effect of Eligibility to the Aide au Mérite in Bac Year on Enrolling in a Selective Masters Degree

	First order			Second order		
	(1)	(2)	(3)	(4)	(5)	(6)
<i>Panel A: MSE-optimal bandwidth</i>						
Eligibility	-0.002	-0.011	-0.004	0.003	-0.028	-0.022
Robust 95% CI	[-0.02, 0.02]	[-0.05, 0.01]	[-0.03, 0.02]	[-0.03, 0.04]	[-0.09, 0.02]	[-0.08, 0.02]
Robust p-value	0.674	0.264	0.554	0.740	0.182	0.275
# obs. left	33,034	28,440	38,866	44,392	49,542	49,542
# obs. right	34,275	23,236	26,809	37,676	29,168	29,168
Bandwidth	(15.11, 16.89)	(15.12, 16.88)	(14.94, 17.06)	(14.93, 17.07)	(14.8, 17.2)	(14.8, 17.2)
<i>Panel B: (15, 17) bandwidth</i>						
Eligibility	-0.003	-0.009	-0.005	0.005	-0.039	-0.031
Conventional 95% CI	[-0.02, 0.01]	[-0.03, 0.01]	[-0.03, 0.02]	[-0.03, 0.04]	[-0.1, 0.03]	[-0.1, 0.03]
Conventional p-value	0.681	0.426	0.637	0.739	0.253	0.337
# obs. left	39,274	35,234	35,234	39,274	35,234	35,234
# obs. right	35,879	25,357	25,357	35,879	25,357	25,357
Bandwidth	(15, 17)	(15, 17)	(15, 17)	(15, 17)	(15, 17)	(15, 17)
Poly. order	1	1	1	2	2	2
Donut		✓	✓		✓	✓
Controls			✓			✓
Mean [15.5, 15.7)	0.227	0.227	0.227	0.227	0.227	0.227

Notes: This table reports estimates of the discontinuity in the probability of enrolling in a selective masters degree at the aide au mérite eligibility threshold (16/20 Bac grade). Selective masters degrees are defined as degrees for which the final year of study is 4 or 5 and is delivered by an engineering, business or other private school. Panel A reports estimates using the MSE-optimal bandwidth (obtained from the *rdrobust* R package). Panel B reports estimates using the (15, 17) bandwidth. For MSE-optimal bandwidth estimates, the reported confidence intervals correspond to associated robust 95% confidence intervals, while they correspond to associated conventional 95% confidence intervals for (15, 17) bandwidth estimates. For each panel, the first row reports the discontinuity estimate in the outcome at the 16/20 Bac grade threshold. Columns (1)-(3) report estimates obtained from local linear regressions (triangular kernel) while columns (4)-(6) report estimates from regressions including a quadratic of the running variable. Columns (1) and (4) report estimates on the full sample, columns (2) and (5) estimates excluding students with Bac grades in [15.7, 16.05] ("Donut"), and columns (3) and (6) estimates excluding students with Bac grades in [15.7, 16.05] and including control variables (gender, age, SES, Bac track, and Bac cohort). See Table C.1's notes for an example of how to read the estimates. Statistical significance is computed based on the relevant p-value and ***, **, and * indicate significance at 1, 5, and 10%, respectively.

Table C.15: Effect of Eligibility to aide au mérite on First Graduate Degree Quality

	<i>First order</i>			<i>Second order</i>		
	(1)	(2)	(3)	(4)	(5)	(6)
<i>Panel A: MSE-optimal bandwidth</i>						
Eligibility	-0.17**	0.05	0.055	-0.154**	0.031	0.033
Robust 95% CI	[-0.36, -0.02]	[-0.11, 0.18]	[-0.09, 0.18]	[-0.32, -0.02]	[-0.18, 0.21]	[-0.15, 0.19]
Robust p-value	0.029	0.622	0.511	0.024	0.877	0.784
# obs. left	6,192	15,163	16,205	24,721	35,037	38,434
# obs. right	16,799	16,160	16,640	26,891	23,491	24,427
Bandwidth	(15.57, 16.43)	(15.21, 16.79)	(15.19, 16.81)	(15.02, 16.98)	(14.71, 17.29)	(14.64, 17.36)
<i>Panel B: (15, 17) bandwidth</i>						
Eligibility	-0.054*	0.058	0.053	-0.149**	0.017	0.042
Conventional 95% CI	[-0.12, 0.01]	[-0.03, 0.15]	[-0.03, 0.14]	[-0.28, -0.02]	[-0.25, 0.28]	[-0.21, 0.29]
Conventional p-value	0.089	0.195	0.210	0.022	0.901	0.744
# obs. left	25,730	23,059	23,059	25,730	23,059	23,059
# obs. right	26,891	19,286	19,286	26,891	19,286	19,286
Bandwidth	(15, 17)	(15, 17)	(15, 17)	(15, 17)	(15, 17)	(15, 17)
Poly. order	1	1	1	2	2	2
Donut		✓	✓		✓	✓
Controls			✓			✓
Mean [15.5, 15.7)	13.719	13.719	13.719	13.719	13.719	13.719

Notes: This table reports estimates of the discontinuity in (first) graduate degree quality at the aide au mérite eligibility threshold (16/20 Bac grade). Degree quality is measured as the median Bac grade among contemporaneously enrolled students. Panel A reports estimates using the MSE-optimal bandwidth (obtained from the *rdrobust* R package). Panel B reports estimates using the (15, 17) bandwidth. For MSE-optimal bandwidth estimates, the reported confidence intervals correspond to associated robust 95% confidence intervals, while they correspond to associated conventional 95% confidence intervals for (15, 17) bandwidth estimates. For each panel, the first row reports the discontinuity estimate in the outcome at the 16/20 Bac grade threshold. Columns (1)-(3) report estimates obtained from local linear regressions (triangular kernel) while columns (4)-(6) report estimates from regressions including a quadratic of the running variable. Columns (1) and (4) report estimates on the full sample, columns (2) and (5) estimates excluding students with Bac grades in [15.7, 16.05] ("Donut"), and columns (3) and (6) estimates excluding students with Bac grades in [15.7, 16.05] and including control variables (gender, age, SES, Bac track, and Bac cohort). See Table C.1's notes for an example of how to read the estimates. Statistical significance is computed based on the relevant p-value and ***, **, and * indicate significance at 1, 5, and 10%, respectively.

Table C.16: Effect of Eligibility to the Aide au Mérite in Bac Year on Enrolling in Paris (Urban Unit) in Bac Year

	<i>First order</i>			<i>Second order</i>		
	(1)	(2)	(3)	(4)	(5)	(6)
<i>Panel A: MSE-optimal bandwidth</i>						
Eligibility	0.093***	0.03*	0.029*	0.092***	0.024	0.022
Robust 95% CI	[0.07, 0.12]	[-0.01, 0.07]	[-0.01, 0.07]	[0.07, 0.12]	[-0.04, 0.09]	[-0.04, 0.08]
Robust p-value	0.000	0.087	0.088	0.000	0.514	0.522
# obs. left	10,286	15,810	15,810	35,884	33,660	35,234
# obs. right	23,659	18,367	18,367	34,816	25,349	25,357
Bandwidth	(15.55, 16.45)	(15.35, 16.65)	(15.35, 16.65)	(15.06, 16.94)	(15.02, 16.98)	(15.01, 16.99)
<i>Panel B: (15, 17) bandwidth</i>						
Eligibility	0.049***	0.02**	0.018**	0.086***	0.023	0.022
Conventional 95% CI	[0.04, 0.06]	[0, 0.04]	[0, 0.04]	[0.06, 0.11]	[-0.03, 0.08]	[-0.03, 0.07]
Conventional p-value	0.000	0.029	0.043	0.000	0.395	0.410
# obs. left	39,274	35,234	35,234	39,274	35,234	35,234
# obs. right	35,879	25,357	25,357	35,879	25,357	25,357
Bandwidth	(15, 17)	(15, 17)	(15, 17)	(15, 17)	(15, 17)	(15, 17)
Poly. order	1	1	1	2	2	2
Donut		✓	✓		✓	✓
Controls			✓			✓
Mean [15.5, 15.7)	0.129	0.129	0.129	0.129	0.129	0.129

Notes: This table reports estimates of the discontinuity in the probability of enrolling in Bac year a higher education institution located in the Paris (urban unit) at the aide au mérite eligibility threshold (16/20 Bac grade). Panel A reports estimates using the MSE-optimal bandwidth (obtained from the *rdrobust* R package). Panel B reports estimates using the (15, 17) bandwidth. For MSE-optimal bandwidth estimates, the reported confidence intervals correspond to associated robust 95% confidence intervals, while they correspond to associated conventional 95% confidence intervals for (15, 17) bandwidth estimates. For each panel, the first row reports the discontinuity estimate in the outcome at the 16/20 Bac grade threshold. Columns (1)-(3) report estimates obtained from local linear regressions (triangular kernel) while columns (4)-(6) report estimates from regressions including a quadratic of the running variable. Columns (1) and (4) report estimates on the full sample, columns (2) and (5) estimates excluding students with Bac grades in [15.7, 16.05] ("Donut"), and columns (3) and (6) estimates excluding students with Bac grades in [15.7, 16.05] and including control variables (gender, age, SES, Bac track, and Bac cohort). See Table C.1's notes for an example of how to read the estimates. Statistical significance is computed based on the relevant p-value and ***, **, and * indicate significance at 1, 5, and 10%, respectively.

Table C.17: Effect of Eligibility to the Aide au Mérite in Bac Year on Enrolling in Paris, Marseille or Lyon (Urban Units) in Bac Year

	<i>First order</i>			<i>Second order</i>		
	(1)	(2)	(3)	(4)	(5)	(6)
<i>Panel A: MSE-optimal bandwidth</i>						
Eligibility	0.077***	0.04**	0.04**	0.086***	0.042	0.039
Robust 95% CI	[0.06, 0.1]	[0, 0.09]	[0, 0.09]	[0.06, 0.12]	[-0.03, 0.11]	[-0.03, 0.11]
Robust p-value	0.000	0.035	0.032	0.000	0.283	0.292
# obs. left	19,850	17,931	17,177	42,579	36,851	36,991
# obs. right	28,889	19,136	19,022	36,883	25,853	26,327
Bandwidth	(15.35, 16.65)	(15.31, 16.69)	(15.31, 16.69)	(14.96, 17.04)	(14.99, 17.01)	(14.98, 17.02)
<i>Panel B: (15, 17) bandwidth</i>						
Eligibility	0.053***	0.028***	0.027**	0.09***	0.044	0.044
Conventional 95% CI	[0.04, 0.07]	[0.01, 0.05]	[0.01, 0.05]	[0.06, 0.12]	[-0.02, 0.11]	[-0.02, 0.11]
Conventional p-value	0.000	0.009	0.012	0.000	0.180	0.178
# obs. left	39,274	35,234	35,234	39,274	35,234	35,234
# obs. right	35,879	25,357	25,357	35,879	25,357	25,357
Bandwidth	(15, 17)	(15, 17)	(15, 17)	(15, 17)	(15, 17)	(15, 17)
Poly. order	1	1	1	2	2	2
Donut		✓	✓		✓	✓
Controls			✓			✓
Mean [15.5, 15.7)	0.207	0.207	0.207	0.207	0.207	0.207

Notes: This table reports estimates of the discontinuity in the probability of enrolling in Bac year a higher education institution located in the Paris, Marseille or Lyon (urban units) at the aide au mérite eligibility threshold (16/20 Bac grade). Panel A reports estimates using the MSE-optimal bandwidth (obtained from the *rdrobust* R package). Panel B reports estimates using the (15, 17) bandwidth. For MSE-optimal bandwidth estimates, the reported confidence intervals correspond to associated robust 95% confidence intervals, while they correspond to associated conventional 95% confidence intervals for (15, 17) bandwidth estimates. For each panel, the first row reports the discontinuity estimate in the outcome at the 16/20 Bac grade threshold. Columns (1)-(3) report estimates obtained from local linear regressions (triangular kernel) while columns (4)-(6) report estimates from regressions including a quadratic of the running variable. Columns (1) and (4) report estimates on the full sample, columns (2) and (5) estimates excluding students with Bac grades in [15.7, 16.05] ("Donut"), and columns (3) and (6) estimates excluding students with Bac grades in [15.7, 16.05] and including control variables (gender, age, SES, Bac track, and Bac cohort). See Table C.1's notes for an example of how to read the estimates. Statistical significance is computed based on the relevant p-value and ***, **, and * indicate significance at 1, 5, and 10%, respectively.

Table C.18: Effect of Eligibility to the Aide au Mérite in Bac Year on Enrolling in a City with Over 200k Inhabitants in Bac Year

	<i>First order</i>			<i>Second order</i>		
	(1)	(2)	(3)	(4)	(5)	(6)
<i>Panel A: MSE-optimal bandwidth</i>						
Eligibility	0.033***	0.033	0.027*	0.046***	0.035	0.031
Robust 95% CI	[0.01, 0.06]	[-0.01, 0.07]	[0, 0.06]	[0.01, 0.08]	[-0.02, 0.09]	[-0.02, 0.08]
Robust p-value	0.003	0.102	0.070	0.005	0.192	0.224
# obs. left	33,034	25,535	33,599	49,538	56,892	64,603
# obs. right	34,275	22,272	24,826	38,803	30,778	32,212
Bandwidth	(15.1, 16.9)	(15.17, 16.83)	(15.04, 16.96)	(14.85, 17.15)	(14.69, 17.31)	(14.6, 17.4)
<i>Panel B: (15, 17) bandwidth</i>						
Eligibility	0.03***	0.029**	0.027**	0.059***	0.061	0.059
Conventional 95% CI	[0.01, 0.05]	[0, 0.05]	[0, 0.05]	[0.02, 0.09]	[-0.01, 0.14]	[-0.01, 0.13]
Conventional p-value	0.001	0.021	0.027	0.001	0.111	0.117
# obs. left	39,274	35,234	35,234	39,274	35,234	35,234
# obs. right	35,879	25,357	25,357	35,879	25,357	25,357
Bandwidth	(15, 17)	(15, 17)	(15, 17)	(15, 17)	(15, 17)	(15, 17)
Poly. order	1	1	1	2	2	2
Donut		✓	✓		✓	✓
Controls			✓			✓
Mean [15.5, 15.7)	0.351	0.351	0.351	0.351	0.351	0.351

Notes: This table reports estimates of the discontinuity in the probability of enrolling in Bac year a higher education institution located in a city with over 200,000 inhabitants at the aide au mérite eligibility threshold (16/20 Bac grade). The cities with over 200,000 inhabitants are: Paris, Marseille, Lyon, Toulouse, Nice, Nantes, Montpellier, Strasbourg, Bordeaux, Lille, and Rennes. Panel A reports estimates using the MSE-optimal bandwidth (obtained from the *rdrobust* R package). Panel B reports estimates using the (15, 17) bandwidth. For MSE-optimal bandwidth estimates, the reported confidence intervals correspond to associated robust 95% confidence intervals, while they correspond to associated conventional 95% confidence intervals for (15, 17) bandwidth estimates. For each panel, the first row reports the discontinuity estimate in the outcome at the 16/20 Bac grade threshold. Columns (1)-(3) report estimates obtained from local linear regressions (triangular kernel) while columns (4)-(6) report estimates from regressions including a quadratic of the running variable. Columns (1) and (4) report estimates on the full sample, columns (2) and (5) estimates excluding students with Bac grades in [15.7, 16.05] (“Donut”), and columns (3) and (6) estimates excluding students with Bac grades in [15.7, 16.05] and including control variables (gender, age, SES, Bac track, and Bac cohort). See Table C.1’s notes for an example of how to read the estimates. Statistical significance is computed based on the relevant p-value and ***, **, and * indicate significance at 1, 5, and 10%, respectively.

Table C.19: Effect of Eligibility to the Aide au Mérite in Bac Year on Enrolling in Paris (Urban Unit) in Bac Year + 1

	First order			Second order		
	(1)	(2)	(3)	(4)	(5)	(6)
<i>Panel A: MSE-optimal bandwidth</i>						
Eligibility	0.094***	0.04*	0.039*	0.08***	0.022	0.02
Robust 95% CI	[0.08, 0.12]	[-0.01, 0.09]	[0, 0.09]	[0.06, 0.1]	[-0.04, 0.08]	[-0.04, 0.08]
Robust p-value	0.000	0.082	0.079	0.000	0.499	0.514
# obs. left	10,619	13,831	13,831	46,416	36,894	38,455
# obs. right	24,231	17,209	17,209	38,233	25,853	26,332
Bandwidth	(15.53, 16.47)	(15.39, 16.61)	(15.38, 16.62)	(14.9, 17.1)	(14.99, 17.01)	(14.97, 17.03)
<i>Panel B: (15, 17) bandwidth</i>						
Eligibility	0.049***	0.023***	0.022**	0.089***	0.024	0.024
Conventional 95% CI	[0.04, 0.06]	[0.01, 0.04]	[0, 0.04]	[0.07, 0.11]	[-0.03, 0.08]	[-0.03, 0.08]
Conventional p-value	0.000	0.010	0.015	0.000	0.365	0.372
# obs. left	39,274	35,234	35,234	39,274	35,234	35,234
# obs. right	35,879	25,357	25,357	35,879	25,357	25,357
Bandwidth	(15, 17)	(15, 17)	(15, 17)	(15, 17)	(15, 17)	(15, 17)
Poly. order	1	1	1	2	2	2
Donut		✓	✓		✓	✓
Controls			✓			✓
Mean [15.5, 15.7)	0.128	0.128	0.128	0.128	0.128	0.128

Notes: This table reports estimates of the discontinuity in the probability of enrolling in Bac year + 1 in a higher education institution located in the Paris (urban unit) at the aide au mérite eligibility threshold (16/20 Bac grade). Panel A reports estimates using the MSE-optimal bandwidth (obtained from the *rdrobust* R package). Panel B reports estimates using the (15, 17) bandwidth. For MSE-optimal bandwidth estimates, the reported confidence intervals correspond to associated robust 95% confidence intervals, while they correspond to associated conventional 95% confidence intervals for (15, 17) bandwidth estimates. For each panel, the first row reports the discontinuity estimate in the outcome at the 16/20 Bac grade threshold. Columns (1)-(3) report estimates obtained from local linear regressions (triangular kernel) while columns (4)-(6) report estimates from regressions including a quadratic of the running variable. Columns (1) and (4) report estimates on the full sample, columns (2) and (5) estimates excluding students with Bac grades in [15.7, 16.05] (“Donut”), and columns (3) and (6) estimates excluding students with Bac grades in [15.7, 16.05] and including control variables (gender, age, SES, Bac track, and Bac cohort). See Table C.1’s notes for an example of how to read the estimates. Statistical significance is computed based on the relevant p-value and ***, **, and * indicate significance at 1, 5, and 10%, respectively.

Table C.20: Effect of Eligibility to the Aide au Mérite in Bac Year on Enrolling in Paris, Marseille or Lyon (Urban Units) in Year Bac + 1

	<i>First order</i>			<i>Second order</i>		
	(1)	(2)	(3)	(4)	(5)	(6)
<i>Panel A: MSE-optimal bandwidth</i>						
Eligibility	0.077***	0.028	0.028	0.075***	0.022	0.024
Robust 95% CI	[0.06, 0.1]	[-0.01, 0.08]	[-0.01, 0.07]	[0.06, 0.1]	[-0.04, 0.09]	[-0.03, 0.09]
Robust p-value	0.000	0.124	0.120	0.000	0.458	0.398
# obs. left	21,217	17,931	18,374	60,206	42,043	43,505
# obs. right	29,544	19,136	19,284	41,228	27,307	27,783
Bandwidth	(15.32, 16.68)	(15.3, 16.7)	(15.3, 16.7)	(14.71, 17.29)	(14.9, 17.1)	(14.88, 17.12)
<i>Panel B: (15, 17) bandwidth</i>						
Eligibility	0.052***	0.023**	0.022**	0.096***	0.028	0.028
Conventional 95% CI	[0.04, 0.07]	[0, 0.04]	[0, 0.04]	[0.07, 0.12]	[-0.04, 0.09]	[-0.04, 0.09]
Conventional p-value	0.000	0.038	0.047	0.000	0.394	0.384
# obs. left	39,274	35,234	35,234	39,274	35,234	35,234
# obs. right	35,879	25,357	25,357	35,879	25,357	25,357
Bandwidth	(15, 17)	(15, 17)	(15, 17)	(15, 17)	(15, 17)	(15, 17)
Poly. order	1	1	1	2	2	2
Donut		✓	✓		✓	✓
Controls			✓			✓
Mean [15.5, 15.7)	0.21	0.21	0.21	0.21	0.21	0.21

Notes: This table reports estimates of the discontinuity in the probability of enrolling in Bac year + 1 in a higher education institution located in the Paris, Marseille or Lyon (urban units) at the aide au mérite eligibility threshold (16/20 Bac grade). Panel A reports estimates using the MSE-optimal bandwidth (obtained from the *rdrobust* R package). Panel B reports estimates using the (15, 17) bandwidth. For MSE-optimal bandwidth estimates, the reported confidence intervals correspond to associated robust 95% confidence intervals, while they correspond to associated conventional 95% confidence intervals for (15, 17) bandwidth estimates. For each panel, the first row reports the discontinuity estimate in the outcome at the 16/20 Bac grade threshold. Columns (1)-(3) report estimates obtained from local linear regressions (triangular kernel) while columns (4)-(6) report estimates from regressions including a quadratic of the running variable. Columns (1) and (4) report estimates on the full sample, columns (2) and (5) estimates excluding students with Bac grades in [15.7, 16.05] (“Donut”), and columns (3) and (6) estimates excluding students with Bac grades in [15.7, 16.05] and including control variables (gender, age, SES, Bac track, and Bac cohort). See Table C.1’s notes for an example of how to read the estimates. Statistical significance is computed based on the relevant p-value and ***, **, and * indicate significance at 1, 5, and 10%, respectively.

Table C.21: Effect of Eligibility to the Aide au Mérite in Bac Year on Enrolling in City with Over 200k Inhabitants in Bac Year + 1

	First order			Second order		
	(1)	(2)	(3)	(4)	(5)	(6)
<i>Panel A: MSE-optimal bandwidth</i>						
Eligibility	0.039***	0.059**	0.048**	0.049***	0.048	0.046*
Robust 95% CI	[0.02, 0.07]	[0.01, 0.11]	[0.01, 0.09]	[0.02, 0.09]	[-0.02, 0.12]	[-0.01, 0.1]
Robust p-value	0.001	0.010	0.011	0.003	0.141	0.079
# obs. left	29,575	19,442	22,674	49,538	45,498	54,423
# obs. right	32,794	19,862	21,461	38,803	28,281	30,397
Bandwidth	(15.16, 16.84)	(15.27, 16.73)	(15.21, 16.79)	(14.86, 17.14)	(14.86, 17.14)	(14.72, 17.28)
<i>Panel B: (15, 17) bandwidth</i>						
Eligibility	0.033***	0.038***	0.036***	0.062***	0.08**	0.078**
Conventional 95% CI	[0.02, 0.05]	[0.01, 0.06]	[0.01, 0.06]	[0.03, 0.1]	[0.01, 0.15]	[0.01, 0.15]
Conventional p-value	0.000	0.003	0.003	0.000	0.035	0.035
# obs. left	39,274	35,234	35,234	39,274	35,234	35,234
# obs. right	35,879	25,357	25,357	35,879	25,357	25,357
Bandwidth	(15, 17)	(15, 17)	(15, 17)	(15, 17)	(15, 17)	(15, 17)
Poly. order	1	1	1	2	2	2
Donut		✓	✓		✓	✓
Controls			✓			✓
Mean [15.5, 15.7)	0.337	0.337	0.337	0.337	0.337	0.337

Notes: This table reports estimates of the discontinuity in the probability of enrolling in Bac year + 1 in a higher education institution located in a city with over 200,000 inhabitants at the aide au mérite eligibility threshold (16/20 Bac grade). The cities with over 200,000 inhabitants are: Paris, Marseille, Lyon, Toulouse, Nice, Nantes, Montpellier, Strasbourg, Bordeaux, Lille, and Rennes. Panel A reports estimates using the MSE-optimal bandwidth (obtained from the *rdrobust* R package). Panel B reports estimates using the (15, 17) bandwidth. For MSE-optimal bandwidth estimates, the reported confidence intervals correspond to associated robust 95% confidence intervals, while they correspond to associated conventional 95% confidence intervals for (15, 17) bandwidth estimates. For each panel, the first row reports the discontinuity estimate in the outcome at the 16/20 Bac grade threshold. Columns (1)-(3) report estimates obtained from local linear regressions (triangular kernel) while columns (4)-(6) report estimates from regressions including a quadratic of the running variable. Columns (1) and (4) report estimates on the full sample, columns (2) and (5) estimates excluding students with Bac grades in [15.7, 16.05] (“Donut”), and columns (3) and (6) estimates excluding students with Bac grades in [15.7, 16.05] and including control variables (gender, age, SES, Bac track, and Bac cohort). See Table C.1’s notes for an example of how to read the estimates. Statistical significance is computed based on the relevant p-value and ***, **, and * indicate significance at 1, 5, and 10%, respectively.

Institut d'études politiques de Paris
ÉCOLE DOCTORALE DE SCIENCES PO
Programme doctoral en économie
Département d'économie
Doctorat en sciences économiques

Mobilité sociale intergénérationnelle

Mesure, mécanismes, et politiques publiques

Gustave KENEDI

Thèse dirigée par

Pierre-Philippe COMBES, Directeur de recherche, CNRS, Sciences Po

soutenance prévue le 21 novembre, 2023

Jury

Pierre CAHUC, Professeur des Universités, Sciences Po

Gabrielle FACK, Professeure des Universités, Université Paris Dauphine - PSL
(*rapportrice*)

Lindsey MACMILLAN, Professor of Economics, UCL (*rapportrice*)

Roland RATHELOT, Associate Professor of Economics, ENSAE-CREST

Résumé en Français

“There is no extravagance more prejudicial to the growth of national wealth than that wasteful negligence which allows genius that happens to be born of lowly parentage to expend itself in lowly work.”

– Alfred Marshall, *Principles of Economics* (1890)

“The extent to which natural capacities develop and reach fruition is affected by all kinds of social conditions and class attitudes. Even the willingness to make an effort, to try, and so to be deserving in the ordinary sense is itself dependent upon happy family and social circumstances. It is impossible in practice to secure equal chances of achievement and culture for those similarly endowed, and therefore we may want to adopt a principle which recognizes this fact and also mitigates the arbitrary effects of the natural lottery itself.”

– John Rawls, *A Theory of Justice* (1971)

DANS quelle mesure les circonstances de l’enfance, telles que le revenu des parents, le quartier de résidence et les enseignants, influencent-elles les trajectoires de vie des individus ? Aux États-Unis, seulement 7,5% des enfants nés au début des années 1980 de parents appartenant au bas de la distribution des revenus parviennent à atteindre le top 20% en tant qu’adultes (Chetty et al., 2014a). Dans une société où il n’y aurait aucune corrélation entre les revenus des parents et ceux de leurs enfants, cette probabilité serait de 20%, car le quintile de revenu d’origine ne serait pas lié aux revenus futurs. Quelles explications peuvent être avancées pour expliquer cette persistance des inégalités de revenus à travers les générations ? Quelles politiques pourraient contribuer à atténuer ces inégalités intergénérationnelles ? Comment d’autres pays se situent-ils par rapport aux États-Unis à cet égard ? En particulier, la France, caractérisée par une moindre inégalité des revenus et un enseignement supérieur relativement abordable, offre-t-elle une plus grande mobilité intergénérationnelle ?

Cette thèse se situe au cœur de ces interrogations, cherchant modestement à approfondir notre compréhension des réponses potentielles. Dans le chapitre 1, en collaboration avec Louis Sirugue, j'explore l'étendue de la persistance des revenus entre générations en France, tout en comparant ces résultats à d'autres économies avancées. Le chapitre 2 se penche sur l'un des mécanismes sous-jacents à l'immobilité intergénérationnelle, en analysant comment les élèves prennent leurs décisions concernant leur choix de formation dans l'enseignement supérieur, notamment en examinant comment les trajectoires des anciens élèves de leur école influencent ces choix. Le chapitre 3 évalue la possibilité de réduire les disparités d'inscription et de réussite dans l'enseignement supérieur entre les étudiant.e.s de milieux modestes et à potentiel scolaire élevé par rapport à leurs pair.e.s issus de milieux aisés, en leur fournissant un soutien financier supplémentaire s'il.elle.s s'inscrivent dans le supérieur.

Le premier chapitre de cette thèse présente de nouvelles estimations de la mobilité intergénérationnelle des revenus en France, en se concentrant sur les enfants nés dans les années 1970. Étonnamment, la connaissance de la mobilité intergénérationnelle de revenus en France demeure lacunaire, surtout compte tenu du fait que le pays est associé à la renommée de Bourdieu. Les études existantes sont basées sur de petits échantillons, des revenus autodéclarés, ont principalement examiné les relations père-fils, et ne sont plus à jour concernant les mesures de mobilité intergénérationnelle mise en avant dans la littérature de cette dernière décennie (Lefranc and Trannoy, 2005; Lefranc, 2018). Une nouvelle source de données administratives riche, l'*Échantillon Démographique Permanent*, a été mise à disposition des chercheur.euse.s. Dans ce chapitre, nous exploitons cette base de données pour estimer la mobilité intergénérationnelle sur un échantillon beaucoup plus vaste, en bénéficiant d'informations plus précises sur les revenus individuels. Notre approche novatrice consiste à regrouper les données concernant les fils et les filles, à intégrer les mères (et les filles) dans l'analyse en définissant le revenu au niveau du *ménage*, et à utiliser des mesures de mobilité intergénérationnelle plus récentes, notamment la corrélation rang-rang. Bien que les revenus des parents ne soient toujours pas observés, nous exploitons les informations abondantes à leur sujet, telles que leur niveau d'éducation, leur profession, leur lieu de résidence, etc., pour prédire leurs revenus en utilisant la méthode des moindres carrés en deux étapes à deux échantillons (TSTSLs).

Nous constatons que la France se caractérise par une forte persistance des revenus à travers les générations par rapport aux autres économies avancées. Seuls 9,7% des enfants issus des 20% des familles les plus pauvres atteignent le top 20% en tant qu'adultes, soit presque 4 fois moins que les enfants nés de parents du haut de la distribution (38,4%).

En comparaison, la probabilité pour un enfant né d'une famille du bas de la distribution des revenus d'atteindre le top 20% à l'âge adulte est de 7,5% aux États-Unis (Chetty et al., 2014a) et de 12,3% en Australie (Deutscher and Mazumder, 2020). Nous documentons également de très fortes variations spatiales de la mobilité intergénérationnelle entre départements. Elles semblent être principalement liées au taux de chômage. Enfin, nous constatons d'importants gains en termes de mobilité géographique. En particulier, le rang de revenu attendu des individus issus du bas de la distribution des revenus des parents qui ont déménagé vers des départements à revenus élevés est à peu près le même que le rang de revenu attendu des individus du 75^{ème} percentile qui sont restés dans leur département d'enfance.

Les résultats de cette étude soulèvent de nouvelles questions stimulantes. Pourquoi la mobilité intergénérationnelle est-elle si limitée en France ? Plus généralement, quels facteurs sous-tendent la mobilité intergénérationnelle ? Pourquoi les enfants issus de familles à revenu élevé ont-ils tendance à obtenir des revenus similaires ? Pourquoi certaines nations connaissent-elles une mobilité intergénérationnelle plus significative que d'autres ? De multiples éléments peuvent contribuer à éclaircir les causes profondes de la mobilité intergénérationnelle. Bowles (1973) a classé ces déterminants en trois catégories : (i) les inégalités dans les opportunités éducatives, (ii) les différences dans les aspirations, les traits de personnalité et d'autres caractéristiques culturelles liées à l'environnement familial, et (iii) l'héritabilité des capacités intellectuelles. Les chapitres 2, réalisé en collaboration avec Nagui Bechichi, et 3 contribuent à notre compréhension des mécanismes qui se situent au croisement des inégalités dans les opportunités éducatives (notamment supérieures) et des différences dans les aspirations liées au contexte familial.

Le deuxième chapitre explore l'impact de la trajectoire éducative des camarades de lycée de la cohorte précédente sur les choix d'enseignement supérieur des étudiant.e.s. La décision de postuler à l'université et de choisir la bonne formation est complexe. Cependant, cette décision revêt une importance considérable, car l'obtention d'un diplôme d'enseignement supérieur offre l'un des rendements les plus élevés, bien que ces rendements varient en fonction des disciplines et, dans une certaine mesure, des établissements (Altonji et al., 2012; Hastings et al., 2013; Kirkeboen et al., 2016; Aucejo et al., 2022; Black et al., 2023; Chetty et al., 2023). Les étudiant.e.s ne sont pas tou.te.s également préparé.e.s à prendre cette décision, en raison de disparités dans l'information disponible sur les rendements de l'enseignement supérieur, de différences dans la connaissance des diverses institutions et filières, des aspirations influencées par le milieu familial, ou encore des ressources financières disponibles.

Des recherches récentes ont mis en évidence l'importance des liens sociaux des étudiant.e.s, y compris leur famille (parents et frères et sœurs) ainsi que leurs liens plus proches (ami.e.s, voisin.e.s, enseignant.e.s) (Aguirre and Matta, 2021; Altmejd et al., 2021; Barrios-Fernández, 2022; Altmejd, 2023). Ces études suggèrent que l'exposition aux formations du supérieur des camarades de lycée peut jouer un rôle significatif dans l'acquisition d'informations sur l'enseignement supérieur. Dans ce chapitre, nous examinons, pour la première fois, dans quelle mesure les choix d'inscription et de filière des élèves sont influencés par les parcours universitaires des élèves de la cohorte précédente du même lycée.

En utilisant une méthode de régression par discontinuité et des données administratives françaises relatives aux candidatures dans l'enseignement supérieur, nous mettons en évidence des effets de contagion (*spillovers*) importants au sein des lycées. Les étudiant.e.s sont nettement plus enclins à présenter des candidatures et à s'inscrire dans une filière universitaire si, l'année précédente, un.e étudiant.e a été marginalement admis.e dans la même filière universitaire et provenait du même lycée, comparativement aux étudiant.e.s des lycées où un.e étudiant.e a été marginalement refusé.e. De manière générale, il.elle.s ont davantage tendance à postuler dans le même établissement du supérieur, mais nous ne constatons aucun impact sur le choix de filière. De plus, l'ampleur de ces effets varie en fonction du lycée, de la formation et des interactions entre les caractéristiques du lycée et de la formation. Les lycées de plus petite taille présentent des effets de contagion plus prononcés entre les cohortes d'étudiant.e.s, tout comme les filières moins sélectives. La distance géographique semble également jouer un rôle, avec les formations très proches et semi-lointaines ayant les effets de contagion les plus marqués en termes de candidatures. Nous constatons que les effets de *role model* des étudiant.e.s expliquent la majeure partie de ces effets de contagion, plutôt que l'influence des enseignant.e.s. De plus, les filles ont nettement plus tendance à postuler dans une formation du supérieur si l'étudiant.e marginalement admis.e de la cohorte précédente est une femme, tandis que l'inverse est vrai pour les étudiants masculins. Ces résultats mettent en lumière l'importance cruciale de l'environnement au lycée des étudiant.e.s dans la formation de leurs choix d'enseignement supérieur.

Le troisième chapitre évalue si une augmentation de l'aide financière dans le supérieur peut réduire les disparités d'inscription observées entre les étudiant.e.s à haut potentiel scolaire issu.e.s de familles défavorisées et leurs pairs provenant de milieux aisés. Dans de nombreux pays, des écarts significatifs subsistent en matière d'inscription, de qualité des établissements fréquentés et de diplomation entre les étudiant.e.s issu.e.s de différents mi-

lieux socio-économiques, même lorsque l'on tient compte de leurs performances scolaires au lycée (Hoxby and Avery, 2013; Crawford et al., 2016; Dynarski et al., 2021; Campbell et al., 2023; Hakimov et al., 2022). Quelles sont les causes de ces importantes disparités ? S'expliquent-elles par un manque de sensibilisation des étudiant.e.s très bon.e.s scolairement issu.e.s de milieux défavorisés quant aux avantages de l'enseignement supérieur, par un manque d'information sur les programmes pertinents, ou par le besoin de ressources financières supplémentaires pour fréquenter des établissements sélectifs ?

Dans ce chapitre, j'évalue l'impact de l'octroi automatique d'une aide financière supplémentaire aux étudiant.e.s très bons scolairement issu.e.s de milieux défavorisés qui s'inscrivent dans l'enseignement supérieur. En utilisant des données administratives exhaustives pour la France et une méthode de régression par discontinuité, je constate que cette politique n'a eu aucun effet sur l'inscription, la persistance, la diplomation ou la performance académique dans l'enseignement supérieur. En outre, il n'y a aucune preuve que cette aide ait incité les étudiant.e.s éligibles à s'inscrire ou à se réorienter vers des filières plus sélective au cours de leurs études. Deux principales conclusions émergent de cette étude. Premièrement, du moins dans le contexte français, un soutien financier supplémentaire aux étudiant.e.s boursier.e.s ayant obtenu la mention Très Bien au Bac, sans aucune autre modification, ne semble pas avoir eu d'impact sur les résultats académiques pertinents. Toutefois, cette politique pourrait avoir eu des effets positifs sur la santé mentale et la réduction des difficultés financières des étudiant.e.s, des variables qui ne sont pas observées dans les données. Deuxièmement, à la lumière de ces constatations et des conclusions de la littérature, il semble exister des complémentarités entre le soutien financier et le niveau académique des étudiant.e.s. Plus précisément, les étudiant.e.s ayant un niveau académique moins élevé au moment de leur admission à l'université semblent être plus sensibles aux conséquences négatives du manque de soutien financier par rapport à leurs homologues plus académiquement préparés.

Ci-dessous, je décris chaque chapitre plus en détail.

Chapitre 1: Mobilité Intergénérationnelle de Revenus en France: Une Analyse Comparative et Géographique

co-écrit avec Louis Sirugue (Paris School of Economics)

Dans quelle mesure le revenu des individus est-il lié à celui de leurs parents ? Cette question suscite un intérêt renouvelé tant dans le grand public que dans le milieu universitaire,

car la montée des inégalités de revenus a soulevé des préoccupations concernant l'égalité des chances. Examiner ce lien est essentiel pour comprendre si les enfants issus de milieux socio-économiques différents ont les mêmes opportunités. Cela a également un impact sur l'efficacité économique, car une forte persistance des revenus d'une génération à l'autre peut refléter une allocation inefficace des talents (les "Einsteins perdus"). La persistance intergénérationnelle a désormais été estimée pour un grand nombre de pays, ouvrant la voie à des comparaisons internationales intéressantes. Cependant, beaucoup de choses restent à savoir pour la France, un pays caractérisé par une inégalité de revenu relativement modeste après impôts et transferts en comparaison internationale et des frais de scolarité dans l'enseignement supérieur comparativement faible.

Les rares études existantes pour la France estiment uniquement la traditionnelle élasticité intergénérationnelle des revenus (IGE), qui mesure l'élasticité du revenu de l'enfant par rapport au revenu des parents, et sont basées sur de petites enquêtes avec des revenus déclarés par les répondants (Lefranc and Trannoy, 2005; Lefranc, 2018). En utilisant un grand échantillon combinant les données du recensement et les déclarations fiscales, nous estimons deux mesures supplémentaires de la mobilité intergénérationnelle : (i) la corrélation rang-rang (RRC), de plus en plus commune dans la littérature, qui correspond à la corrélation entre les rangs percentiles des revenus des enfants et des parents, et (ii) les matrices de transition, qui capturent des schémas de mobilité plus fins le long de la distribution des revenus des parents. Alors que les études précédentes en France utilisaient des revenus du travail déclarés par les répondants, nous nous concentrons sur des mesures de revenu au niveau du ménage. Elles offrent une meilleure représentation des ressources économiques d'une personne et permettent d'inclure les enfants élevés par des mères célibataires. L'intégration de ces améliorations issues de la "nouvelle" littérature sur la mobilité intergénérationnelle nous permet de réaliser une comparaison internationale détaillée pour situer la France par rapport à d'autres économies avancées pour lesquelles des estimations comparables sont disponibles.

De plus, nous examinons les variations spatiales de la mobilité intergénérationnelle dans les 96 départements métropolitains français. De telles analyses infranationales, initiées par Chetty et al. (2014a), aident à éclairer les mécanismes qui peuvent sous-tendre la persistance des revenus d'une génération à l'autre. Ils mettent en évidence que les estimations au niveau national ne fournissent qu'une évaluation incomplète de la mobilité intergénérationnelle d'un pays. Nous utilisons la dimension de panel de nos données pour décrire les schémas de mobilité géographique des individus et étudier la relation entre la mobilité géographique et la mobilité intergénérationnelle. Nous examinons les

rôles distincts du déménagement vers un département à revenu plus élevé par rapport à l'ascension de l'échelle des revenus au sein des départements, en fonction du rang de revenu des parents.

Notre analyse porte sur près de 65 000 enfants nés entre 1972 et 1981, et observés dans l'Échantillon Démographique Permanent (EDP). Ce riche ensemble de données administratives nous permet de mettre en œuvre les contributions discutées ci-dessus et de répondre de manière convaincante aux préoccupations liées aux biais de cycle de vie et d'atténuation (Haider and Solon, 2006; Black and Devereux, 2011; Nybom and Stuhler, 2017). Étant donné que les revenus des parents ne sont pas observés, nous utilisons une estimation à deux étapes en deux échantillons (TSTSLs) qui consiste à prédire les revenus des parents en utilisant d'autres parents issus de la même population mais dont les revenus sont observés (Björklund and Jäntti, 1997). Cette méthode a déjà été utilisée pour estimer l'IGE dans le contexte français (Lefranc and Trannoy, 2005; Lefranc, 2018) ainsi que dans de nombreux autres pays (Jerrim et al., 2016, Tableau A1).

Alors que les études utilisent généralement l'éducation et/ou la profession pour prédire le revenu des parents, nous exploitons la richesse de nos données pour inclure également des caractéristiques démographiques détaillées des parents (nationalité française, pays de naissance, structure du ménage et cohorte de naissance), ainsi que des caractéristiques de la commune de résidence (taux de chômage, part de mères célibataires, part d'étrangers, population et densité de population). Nos résultats sont largement insensibles à l'ensemble des prédicteurs. Le revenu des parents est alors défini comme la moyenne des salaires moyens avant impôts du père et de la mère entre 35 et 45 ans, et le revenu des enfants comme le revenu du ménage avant impôts moyen entre 2010 et 2016 pour la même tranche d'âge. Ces deux définitions des revenus représentent les définitions les plus complètes au niveau du ménage disponibles pour chaque génération.

Exercice de validation TSTSLs. En utilisant le Panel Study of Income Dynamics (PSID) des États-Unis, nous constatons que le TSTSLs sous-estime légèrement les mesures de la persistance intergénérationnelle basées sur les rangs par rapport à ce qui serait obtenu si le revenu des parents était observé (OLS). Le biais à la baisse par rapport à l'estimation OLS pour la RRC varie de 11% lorsque l'éducation est le seul prédicteur, à environ 3-5% une fois que la profession est également incluse. Les estimations TSTSLs infranationales sont également assez proches de leurs homologues OLS, bien qu'elles aient tendance à dévier davantage lorsque le nombre d'observations est faible. Nos résultats soulignent que, dans des contextes comme le nôtre, où le revenu des parents ne peut pas être directe-

ment observé, les mesures de mobilité intergénérationnelle basées sur les rangs obtenues avec le TSTSLs fournissent probablement des bornes inférieures raisonnablement proches des vraies estimations. Ces résultats confirment ceux obtenus dans différents contextes et échantillons par [Cortes-Orihuela et al. \(2022\)](#) et [Jacome et al. \(2023\)](#). Nous constatons que ce raisonnement s'applique également à la matrice de transition.

Résultats nationaux. Notre principal constat est que la France présente une persistance intergénérationnelle du revenu relativement forte par rapport à d'autres pays développés. Notre estimation de l'élasticité intergénérationnelle du revenu du ménage est de 0,527, ce qui suggère qu'en moyenne, une augmentation de 10% du revenu des parents est associée à une augmentation de 5,27% du revenu de l'enfant. En d'autres termes, si les parents gagnent 10% de plus que la moyenne des revenus des parents, l'enfant est censé conserver environ 50% de cet avantage relatif. Cette estimation doit être interprétée avec prudence compte tenu de notre exercice de validation, qui suggère que l'IGE TSTSLs est significativement supérieur à la vraie estimation. En appliquant le facteur de correction que nous trouvons, l'IGE diminue à 0,396.

En passant à la relation rang-rang, nous constatons que l'espérance conditionnelle du rang percentile du revenu de l'enfant par rapport au rang percentile du revenu des parents est linéaire le long de la distribution des revenus des parents, avec des relations plus prononcées aux extrémités. Notre estimation de la corrélation rang-rang est de 0,303, ce qui implique qu'une augmentation de 10 percentiles dans le rang de revenu des parents est associée, en moyenne, à une augmentation de 3,03 percentiles dans le rang de revenu de l'enfant. Cette estimation est d'une ampleur similaire à celle trouvée pour l'Italie (0,3 ; [Acciari et al. \(2022\)](#)), un peu plus petite que pour les États-Unis (0,341 ; [Chetty et al. \(2014b\)](#)), et nettement plus grande que les estimations existantes pour d'autres économies avancées telles que la Suède (0,197 ; [Heidrich \(2017\)](#)), l'Australie (0,215 ; [Deutscher and Mazumder \(2020\)](#)) ou le Canada (0,242 ; [Corak \(2020\)](#)). En appliquant le facteur de correction que nous trouvons dans l'exercice de validation, nous obtenons une RRC de 0,314, ce qui n'affecte pas la position relative de la France.

La persistance intergénérationnelle, telle que capturée par la matrice de transition, est la plus forte aux extrémités de la distribution des revenus des parents : 9,7% des enfants issus des 20% les plus pauvres de la distribution des revenus des parents parviennent aux 20% les plus riches à l'âge adulte. Cette probabilité est presque quatre fois plus élevée pour les enfants nés de parents faisant partie des 20% les plus riches (38,4%). En comparaison, la probabilité pour un enfant né d'une famille faisant partie des 20% les plus pauvres

d'atteindre les 20% les plus riches à l'âge adulte est de 7,5% aux États-Unis (Chetty et al., 2014b) et de 12,3% en Australie (Deutscher and Mazumder, 2020). De plus, la persistance au sommet devient de plus en plus forte à mesure que nous nous rapprochons de l'extrémité droite de la distribution des revenus des parents. Tout comme pour la RRC, l'exercice de validation suggère que ces estimations représentent des bornes supérieures (inférieures) de la mobilité (de la persistance).

Nous montrons que ces résultats sont robustes aux biais potentiels. Tout d'abord, nous évaluons leur sensibilité à la spécification de la prédiction du revenu des parents. En particulier, nous vérifions si le fait de faire varier l'ensemble des prédicteurs ou d'utiliser des méthodes d'estimation non paramétriques influence nos estimations. Les estimations de l'IGE sont surestimées lorsqu'on utilise uniquement l'éducation comme prédicteur, tandis que la RRC et les matrices de transition restent étonnamment stables quel que soit l'ensemble des prédicteurs utilisés. Une prédiction légèrement améliorée grâce à l'utilisation de modèles flexibles n'altère pas de manière significative nos estimations. De plus, nous évaluons la sensibilité de nos estimations aux biais liés au cycle de vie et à l'atténuation en faisant varier les âges auxquels les revenus des enfants et des parents sont mesurés ainsi que le nombre d'observations de revenus des parents utilisées. Nos résultats de base ne semblent pas sous-estimer ni surestimer la mobilité intergénérationnelle en raison de la mesure des revenus des enfants et/ou des parents trop tôt ou trop tard dans le cycle de vie ni en raison de la moyenne des revenus sur un nombre insuffisant d'années.

Résultats infranationaux. Nous constatons d'importantes variations spatiales de la mobilité intergénérationnelle entre les départements, comparables à celles observées entre les pays. Nous définissons la localisation des individus en fonction de leur département de résidence en 1990, lorsqu'ils avaient entre 9 et 18 ans. Les niveaux de mobilité intergénérationnelle les plus élevés se trouvent généralement à l'ouest de la France, tandis que les niveaux les plus bas sont observés dans le Nord et le Sud. Par exemple, l'IGE varie de 0,30 à 0,45 dans les départements de Bretagne (à l'ouest), tandis qu'il va de 0,42 à 0,70 dans les départements des Hauts-de-France (au nord). La distribution des RRC au niveau des départements est plus resserrée que celle des IGE, mais elle présente des schémas spatiaux très similaires.

Nous caractérisons également la mobilité ascendante absolue (AUM) des départements, définie comme le rang de revenu attendu des enfants nés de parents au 25^{ème} percentile, qui est obtenu à partir des valeurs prédites de la régression rang-rang au niveau des

départements (Chetty et al., 2014b). La mobilité ascendante absolue va de 36,8 dans le Pas-de-Calais (au nord) à 54,4 en Haute-Savoie (à l'est). Le département de Paris se distingue en termes d'AUM (49,8) mais présente des niveaux de persistance intergénérationnelle moyens en termes d'IGE (0,51) et de RRC (0,28). La corrélation entre les départements entre l'IGE et la RRC est seulement de 0,65, et de -0,55 avec l'AUM. Cela souligne l'importance d'utiliser une variété de mesures de mobilité intergénérationnelle pour caractériser la persistance des revenus d'un pays d'une génération à l'autre (Deutscher and Mazumder, 2023).

Pour mieux comprendre les sources sous-jacentes à ces variations interdépartementales de la mobilité intergénérationnelle, nous entreprenons une analyse corrélationnelle simple. Nous constatons que la mobilité ascendante absolue présente des relations beaucoup plus fortes avec les caractéristiques des départements en général que l'IGE ou la RRC. Cela suggère que les facteurs qui affectent la mobilité absolue peuvent différer de ceux qui affectent la mobilité relative. La seule caractéristique systématiquement corrélée négativement avec la mobilité intergénérationnelle est le taux de chômage. De manière intrigante, nous ne trouvons aucune preuve d'une "Great Gatsby Curve" intra-française¹ pour l'IGE ou la RRC. Cela contraste avec les résultats d'autres pays (Acciari et al., 2022; Chetty et al., 2014b; Corak, 2020).

Enfin, nous procédons à une analyse descriptive de la relation entre la mobilité intergénérationnelle des revenus et la mobilité géographique. Nous observons d'importants gains de rang de revenu attendu pour les individus qui déménagent, et ces gains diminuent légèrement en fonction du rang de revenu des parents. Par exemple, pour les enfants issus des familles du décile le plus bas de revenu, ceux.celles qui déménagent ont un rang attendu d'environ 5,6 percentiles supérieur à celui de ceux.celles qui restent, tandis que cette différence est d'environ 4,4 percentiles pour les enfants issus des familles du décile le plus élevé. Ces gains s'expliquent en partie par le fait que les personnes qui déménagent s'installent généralement dans des départements à revenu plus élevé à l'âge adulte par rapport à celles qui restent. De plus, les personnes qui déménagent atteignent des rangs locaux dans leur département à l'âge adulte qui diffèrent davantage du rang de leurs parents dans le département de leur enfance. En moyenne, les départements de destination sont caractérisés par des niveaux de revenus plus élevés que les départements d'origine, mais cette tendance est plus prononcée aux extrémités de la distribution des revenus des parents. Cependant, quel que soit le rang de revenu des parents, les gains

¹La "Great Gatsby Curve" fait référence à la corrélation positive entre la persistance intergénérationnelle du revenu (définie par l'IGE) et l'inégalité des revenus (définie par l'indice de Gini) observée entre les pays (Corak, 2013).

de mobilité ascendante absolue associés au déménagement vers un département à revenu plus élevé semblent significatifs et augmentent avec le revenu moyen dans le département de destination. Toutes ces constatations combinent des phénomènes d’auto-sélection et d’effets causaux, et nous laissons le soin de distinguer ces deux canaux pour de futures recherches.

Chapter 2: Effets d’Influence des Camarades de Lycée sur les Choix d’Enseignement Supérieur

co-écrit avec Nagui Bechichi (Paris School of Economics)

Comment les étudiant.e.s prennent-il.elle.s leurs décisions concernant leurs candidatures dans l’enseignement supérieur ? Cette question a suscité une attention considérable en raison des importants avantages associés à l’obtention d’un diplôme, ainsi que des grandes disparités entre filières et établissements. Des travaux récents ont mis en évidence le rôle crucial des déficits d’information (Hoxby and Turner, 2013a; Carrell and Sacerdote, 2017), ainsi que l’influence des liens sociaux des étudiant.e.s, tels que leurs parents (Altmejd, 2023), leurs frères et sœurs (Aguirre and Matta, 2021; Altmejd et al., 2021), voire même leurs voisin.e.s (Barrios-Fernández, 2022). Ces constatations suggèrent que l’exposition aux choix d’enseignement supérieur de leurs pairs peut jouer un rôle significatif dans les décisions des étudiant.e.s. Cependant, nous disposons de peu d’informations sur la manière dont les camarades de lycée de la cohorte précédente influencent ces décisions. Plus spécifiquement, au sein d’un même lycée, comment les candidatures et les choix d’inscription des élèves sont-ils influencés par les parcours universitaires des camarades de lycée de la cohorte précédente ? L’identification de tels effets de manière causale s’avère difficile en raison des différences d’élèves entre les lycées et de l’endogénéité des choix universitaires des élèves à leur lycée d’origine.

Cet article fournit les premières preuves causales sur les effets de contagion *spillovers* au sein d’un même lycée² sur les choix d’enseignement supérieur. En utilisant des données administratives sur les candidatures en France couvrant près de 90% des formations du supérieur entre 2013 et 2017 (Bechichi et al., 2021), nous montrons que les étudiant.e.s sont plus susceptibles de postuler et de s’inscrire dans une formation du supérieur si un.e étudiant.e du même lycée s’est inscrit.e dans cette même filière l’année précédente. Nous

²Techniquement, notre analyse est menée au niveau du lycée x filière, car, en France, les filières dans les lycées sont assez séparées, et les formations du supérieur sont souvent spécifiques aux filières. Pour faciliter la lisibilité, nous utilisons le terme “lycée” pour désigner “lycée x filière”.

constatons également d'importants effets de contagion sur le choix de l'établissement du supérieur en général, mais aucun effet sur la discipline choisie.

Nous identifions les effets de contagion au sein des lycées en exploitant les seuils d'admission générés par la procédure d'admission centralisée en France. Cette allocation garantit que les formations ne peuvent pas anticiper a priori le lycée du dernier étudiant.e admis.e. Ainsi, les lycées autour du seuil d'admission d'une formation sont essentiellement identiques, à l'exception d'avoir un.e étudiant.e classé.e juste au-dessus ou juste en dessous du rang du dernier étudiant.e admis.e dans cette formation. Cela génère des variations quasi aléatoires dans les formations auxquelles les élèves d'un même lycée sont admis et s'inscrivent, ce qui génère à son tour des variations quasi aléatoires dans les filières auxquelles la cohorte suivante d'élèves du même lycée est exposée. Cela nous permet de mettre en œuvre une régression par discontinuité pour estimer les effets de contagion au sein du lycée sur les candidatures et les inscriptions. Alors que les travaux existants ont exploité des seuils comparables générés par des seuils académiques dans les politiques d'admission (par exemple, [Altmejd et al. \(2021\)](#); [Estrada et al. \(2022\)](#)), notre conception est très similaire dans l'esprit, à la différence que nous n'observons que le classement relatif des étudiant.e.s par les formations auxquelles il.elle.s ont postulé. Étant donné que plusieurs étudiant.e.s du même lycée peuvent postuler à la même formation, nous ne conservons que le.la candidat.a le.la mieux classé.e du lycée par la formation, comme dans [Estrada et al. \(2022\)](#).

Nous constatons que les étudiant.e.s suivent les choix d'enseignement supérieur de la cohorte précédente de leur lycée. Ils sont de 7 points de pourcentage (+25% par rapport à la moyenne contrefactuelle) plus susceptibles de candidater dans une formation et de 3 points de pourcentage (+67%) plus susceptibles de s'inscrire dans une formation dans laquelle un.e étudiant.e de la cohorte précédente de leur lycée a été marginalement admis.e et inscrit.e, par rapport aux étudiant.e.s des lycées ayant un.e camarade marginalement refusé.e. Nous trouvons également d'importants impacts sur la marge intensive, c'est-à-dire le nombre de candidatures et d'étudiant.e.s inscrit.e.s : augmentations de 0,24 (+30%) et de 0,05 (+72%) points de pourcentage respectivement.

L'ampleur de ces effets est importante. Comparés aux effets de contagion de formation entre frères et sœurs estimés par [Altmejd et al. \(2021\)](#) pour le Chili, la Croatie et la Suède, nos effets de contagion au sein du lycée sont de 43% à 78% aussi importants que l'impact sur les candidatures qu'ils trouvent, et de 40% à 88% de leurs effets sur les inscriptions. De plus, nous montrons également que les étudiant.e.s sont plus susceptibles de postuler et de s'inscrire dans le même établissement que leurs camarades de lycée de la cohorte

précédente, mais il n’y a pas d’effet de contagion sur le choix de filière. L’ampleur (relative) des effets de contagion au niveau établissement est à peu près similaire à celle des effets de contagion au niveau formation, de 9,6 points de pourcentage (+17%) pour les candidatures et de 10,9 points de pourcentage (+52%) pour les inscriptions. L’absence d’effets de contagion sur les filières pourrait potentiellement s’expliquer par le fait que les étudiant.e.s ont des préférences plus marquées pour ce qu’il.elle.s veulent étudier que pour l’endroit où il.elle.s veulent étudier, ou parce qu’il.elle.s sont plus conscient.e.s des filières existantes. Par conséquent, la composante établissement de l’inscription d’un.e camarade de classe de la cohorte précédente est plus saillante pour eux.elles. Ce résultat est en accord avec [Altmejd et al. \(2021\)](#) et [Aguirre and Matta \(2021\)](#), qui ne trouvent également aucun effet de contagion de filières entre frères et sœurs.

Nous découvrons plusieurs hétérogénéités instructives en ce qui concerne les effets de contagion. Tout d’abord, nous constatons que l’ampleur des effets de contagion est globalement constante sur les quatre années observées. Cela suggère qu’ils ne résultent pas des particularités d’une année donnée, mais plutôt d’un déterminant structurel des choix d’enseignement supérieur des étudiant.e.s. Deuxièmement, en ce qui concerne les caractéristiques des étudiant.e.s, nous constatons que les effets de contagion au sein du lycée sur les candidatures sont d’une ampleur similaire pour les deux sexes, bien que les effets sur les inscriptions soient nettement plus importants pour les garçons. Cela pourrait être dû à des différences dans les types de formations demandées. De plus, et assez étonnamment, nous constatons que les étudiants défavorisés (basé sur la profession du tuteur légal) ne sont que légèrement plus réactifs que leurs pairs très favorisés. Cela est quelque peu inattendu, car a priori, on pourrait penser que les étudiant.e.s très favorisé.e.s sont mieux informé.e.s sur l’enseignement supérieur et donc moins influencé.e.s par la trajectoire d’enseignement supérieur de leurs camarades de lycée de la cohorte précédente. En revanche, les étudiant.e.s défavorisé.e.s ont tendance à être moins au courant du paysage de l’enseignement supérieur ([Hoxby and Turner, 2013b](#)), et l’on pourrait donc s’attendre à ce qu’il.elles soient plus influencé.e.s par leurs pairs.

Troisièmement, l’ampleur des effets de contagion varie en fonction de certaines caractéristiques des lycées. Toutes les filières de lycée présentent des effets de contagion d’une ampleur à peu près similaire, avec des effets légèrement plus importants (en termes de pourcentage) pour la filière littéraire. Ce résultat est assez remarquable car il suggère que l’acquisition d’informations sur les choix d’enseignement supérieur est pertinente dans des contextes très différents. Cela étant dit, les effets de contagion sont plus importants dans les petits lycées (moins de 30 étudiant.e.s) et diminuent en fonction de la taille

du lycée. Cela pourrait être le résultat de plusieurs explications. Les petits lycées peuvent entretenir des relations plus étroites entre les étudiant.e.s et leurs enseignant.e.s, de sorte que les enseignant.e.s sont peut-être mieux informé.e.s des choix d'enseignement supérieur de leurs élèves. En tant que tels, il.elle.s peuvent encourager leurs classes ultérieures à postuler aux mêmes filières que leurs ancien.ne.s élèves. Une autre explication pourrait être que les petits lycées entretiennent de meilleurs liens avec leurs ancien.ne.s élèves grâce, par exemple, à des rencontres annuelles d'ancien.ne.s élèves. Une dernière explication pourrait simplement être que les petits lycées sont situés dans des zones plus rurales où les informations sur les filières d'enseignement supérieur peuvent être plus rares et donc les chocs d'information sont amplifiés beaucoup plus qu'en milieu urbain, où les informations sur l'enseignement supérieur sont plus riches. De plus, nous constatons que les lycées du deuxième et du cinquième quintile dans la distribution du niveau académique des lycées (mesurée par la médiane des notes au Bac de ses élèves) présentent les plus importants effets de contagion en termes de candidatures. Il n'est pas clair exactement ce qui pourrait expliquer ce résultat.

Quatrièmement, nous explorons comment les effets de contagion varient en fonction des caractéristiques des formations du supérieur. Nous constatons que les effets de contagion sont plus importants pour les universités publiques, les IUT et les BTS, mais pas pour les classes préparatoires, qui ont tendance à être assez prestigieuses, ni pour d'autres types de filières. En ligne avec ces résultats, les formations classées dans les 10% les moins sélectives (déterminée par la médiane des notes au Bac des étudiant.e.s inscrit.e.s) présentent les plus importants effets de contagion, et ils diminuent avec la sélectivité. Il n'y a pas d'effets de contagion pour les formations classées dans les 10% les plus sélectives. Cela est quelque peu surprenant, car on pourrait s'attendre à ce que les formations très sélectives soient celles où certains étudiant.e.s n'osent peut-être pas postuler. Cela nous conduit à penser que les étudiant.e.s en apprennent davantage sur les formations du supérieur dont il.elle.s n'avaient pas conscience auparavant, plutôt que d'augmenter leur confiance pour postuler à des filières universitaires prestigieuses.

Enfin, nous évaluons comment l'interaction entre les caractéristiques des lycées et des formations peut façonner les effets de contagion au sein des lycées. Les résultats suggèrent que les formations géographiquement proches (moins de 25 km) et modérément éloignés (entre 50 et 100 km) induisent les plus importants effets de contagion. En termes de marge intensive des candidatures et des inscriptions, ces formations modérément éloignés présentent des effets de contagion significatifs entre les cohortes. De plus, nous constatons que les étudiant.e.s des lycées peu performant académiquement (dans les 25% les moins per-

formants) sont significativement plus susceptibles de suivre un.e camarade de classe marginalement admis.e dans une formation située dans le quartile supérieur de la sélectivité. Cet effet peut être interprété comme relevant les aspirations ou la sensibilisation des étudiant.e.s les plus performants de ces lycées. En revanche, nous constatons de manière intrigante que les effets de contagion sont assez importants pour les lycées les plus performants académiquement (dans les 25% les plus performants) pour lesquels l'étudiant.e marginalement admis.e s'est inscrit.e dans une formation située dans le quartile inférieur de la sélectivité.

Dans la dernière section de l'article, nous explorons deux mécanismes qui pourraient sous-tendre les effets de contagion au sein du lycée : (i) le rôle des enseignant.e.s, et (ii) les effets de modèles (*role model*) d'étudiant.e.s. Bien que ces mécanismes ne soient pas mutuellement exclusifs, ils entraînent des implications de politiques publiques divergentes. Nous évaluons dans quelle mesure les effets de contagion entre les cohortes peuvent être attribués aux enseignant.e.s, par exemple, en influençant les choix de filières de leurs élèves. Étant donné que nous ne disposons pas d'informations directes sur tous les enseignant.e.s des étudiant.e.s, nous explorons ce mécanisme de deux manières complémentaires. Tout d'abord, nous analysons si les étudiant.e.s sont plus enclin.e.s à suivre un.e camarade de classe de la cohorte précédente s'ils partagent le même professeur principal. En France, chaque classe est assignée à un.e professeur.e principal.e responsable des tâches administratives de la classe tout au long de l'année académique, notamment l'aide et la supervision des candidatures dans le supérieur des étudiant.e.s. Deuxièmement, nous évaluons si les étudiant.e.s sont plus susceptibles de suivre un.e camarade de classe de la cohorte précédente s'il.elles sont dans le même numéro de classe (par exemple, classe de terminale S1, classe de terminale S2), ce qui est un indicateur imparfait du partage du même ensemble d'enseignant.e.s que le camarade de classe de la cohorte précédente. Nos résultats indiquent que les étudiant.e.s partageant le même professeur.e principal.e ou le même numéro de classe sont tout aussi enclin.e.s à suivre les choix d'enseignement supérieur de l'étudiant.e de la cohorte précédente marginalement admis.e que les étudiant.e.s ayant des enseignant.e.s différent.e.s ou un.e professeur.e principal.e différent.e. Cela semble suggérer un rôle direct relativement limité des enseignant.e.s, du moins en ce qui concerne l'explication des effets de contagion au sein du lycée que nous mettons en lumière. Cette observation pourrait s'expliquer par le fait que les enseignant.e.s aident leurs élèves en recommandant un large éventail de formations du supérieur plutôt que de se limiter aux filières de leurs ancien.ne.s élèves.

Deuxièmement, nous tentons de distinguer si nos effets de contagion sont plus suscepti-

bles d'être dus à des chocs d'information ou à des effets de rôle modèles d'étudiant.e.s. Pour tester cela, nous évaluons si les effets sont plus importants pour les étudiant.e.s partageant le même sexe ou le même statut socio-économique que l'étudiant.e de la cohorte précédente marginalement admis.e. Nous interprétons ce test comme capturant un effet de rôle modèle d'étudiant.e.s plutôt qu'un effet d'information, car, a priori, le sexe ou le statut socio-économique de l'étudiant.e marginalement admis.e n'affecte pas le contenu informationnel de sa trajectoire d'enseignement supérieur, mais affecte la manière dont cette information est perçue. Nous trouvons de solides preuves en faveur d'effets de rôle modèle d'étudiant.e.s. Les filles sont significativement plus susceptibles de postuler dans des formations lorsque l'étudiant.e de la précédente cohorte marginalement admis.e était une fille (+9%), mais pas lorsqu'il s'agissait d'un garçon (+3%, non significatif), tandis que les garçons sont plus susceptibles de suivre un garçon (+8%), mais pas une fille (+2%, non significatif). De même, les étudiant.e.s défavorisé.e.s sont significativement plus susceptibles de postuler dans une formation lorsque l'étudiant.e de la cohorte précédente marginalement admis.e est également issu.e d'un milieu défavorisé (+13%), mais pas lorsque ce.tte dernier.e provient d'un milieu socio-économique très favorisé (+1%, non significatif). Cependant, les étudiant.e.s très favorisé.e.s sont largement insensibles, quel que soit le statut socio-économique de l'étudiant.e de la cohorte précédente traité.e. Cela est conforme à l'idée qu'il.elle.s ont une meilleure connaissance des filières universitaires ou des préférences plus marquées pour celles-ci.

Chapitre 3 : Étudiant.e.s Défavorisé.e.s, Très Bon.ne.s Scolairement et Aides Financières dans l'Enseignement Supérieur

L'obtention d'un diplôme de l'enseignement supérieur offre l'un des meilleurs rendements qu'un individu puisse réaliser, surtout lorsqu'il.elle fréquente une institution sélective (Bleemer, 2021; Black et al., 2023; Chetty et al., 2023). Cependant, les étudiant.e.s défavorisé.e.s, très bon.ne.s scolairement s'inscrivent dans l'enseignement supérieur à des taux moins élevés que leurs pairs favorisés, et lorsqu'il.elle.s le font, il.elle.s ont tendance à fréquenter des établissements de moins bonne qualité (Hoxby and Avery, 2013; Crawford et al., 2016; Dynarski et al., 2021; Hakimov et al., 2022; Campbell et al., 2022). Ce sous-appariement entraîne d'importantes pertes d'efficacité qui pourraient potentiellement être atténuées par des politiques publiques. Comprendre les facteurs à l'origine de ces écarts est donc essentiel pour concevoir des politiques publiques efficaces. Les étudiant.e.s défavorisé.e.s, très bon.ne.s scolairement sont-il.elle.s moins conscient.e.s des avantages de fréquenter

l'enseignement supérieur, et plus précisément des établissements sélectifs ? Manquent-il.elle.s d'informations sur les formations pertinentes ou n'ont-il.elle.s tout simplement pas la confiance en eux.elles pour postuler ? Ou est-ce qu'il.elle.s ont besoin de ressources financières supplémentaires pour fréquenter ces formations sélectives ? Si les premières raisons prévalent, alors les interventions informatives/motivantes devraient être privilégiées. Cependant, si les contraintes financières sont l'explication dominante, alors un soutien financier ciblé serait la politique préférée.

Dans cet article, j'analyse si une aide financière supplémentaire peut servir de moyen efficace pour inciter les étudiant.e.s défavorisé.e.s, très bon.ne.s scolairement à poursuivre des études supérieures et à s'inscrire dans des établissements de grande qualité, ainsi qu'à persévérer et à obtenir leur diplôme en temps voulu. Plus précisément, j'évalue les effets d'un programme national d'aides financières, l'*aide au mérite*, introduit en 2008 en France, qui accordait automatiquement une allocation supplémentaire de 1 800 euros par an, pendant au plus 3 ans (la durée d'une licence), aux étudiant.e.s éligibles qui s'inscrivaient dans un établissement d'enseignement supérieur. Les seuls critères d'éligibilité à l'aide au mérite étaient que l'étudiant.e (i) soit éligible au programme national de bourses sur critères sociaux, et (ii) obtienne au moins 16 sur 20 (soit dans le top 4,7% des candidat.e.s) à l'examen de fin d'études secondaires français, le *Baccalauréat* (dorénavant le Bac).

La population ciblée d'étudiant.e.s correspond donc très étroitement à la définition des étudiant.e.s défavorisé.e.s, très bon.ne.s scolairement de (Hoxby and Avery, 2013) (top 4% des étudiant.e.s de l'enseignement secondaire américain, et dans le quartile de revenu parental le plus bas). Par définition, l'aide au mérite a été accordée en plus des bourses sur critères sociaux, qui comprenaient une exonération des frais de scolarité et des allocations annuelles pouvant atteindre 5 500 euros pour les étudiant.e.s les plus défavorisé.e.s. En tant que telle, l'aide au mérite représentait au moins une majoration de 40% des allocations mensuelles, une augmentation substantielle du soutien financier.

En utilisant des données administratives sur l'ensemble des étudiant.e.s ayant obtenu le Bac entre 2009 et 2014, j'exploite la discontinuité nette de l'éligibilité à l'aide au mérite au seuil de la note de 16/20 au Bac dans une régression sur discontinuité. Cela me permet d'estimer l'effet causal de l'éligibilité à cette aide financière supplémentaire lors de l'année du Bac sur l'inscription, la qualité du diplôme, la persévérance, l'obtention du diplôme et les performances académiques dans l'enseignement supérieur, ainsi que la mobilité géographique.

Je constate que l'éligibilité à l'aide au mérite lors de l'année du Bac n'a eu aucun effet sur

l'inscription, la persévérance ou l'obtention d'un diplôme de l'enseignement supérieur. Pour la plupart des résultats, je peux exclure des effets aussi faibles que de un à trois points de pourcentage. Dans ce contexte, la marge d'inscription n'est pas particulièrement informative car, conditionnellement à l'éligibilité à une bourse sur critères sociaux, le taux d'inscription autour du seuil de 16 est de 94%. De plus, les étudiant.e.s ne prennent conscience de leur éligibilité à l'aide au mérite qu'en juillet, lorsque les notes au Bac sont publiées, ce qui pourrait limiter l'impact potentiel sur l'inscription. Cependant, comme aux États-Unis, la persévérance dans l'enseignement supérieur est une préoccupation majeure en France. Autour du seuil de 16, moins de trois étudiant.e.s sur quatre éligibles à une bourse sur critères sociaux sont inscrit.e.s en deuxième année à temps, et seulement un peu plus de la moitié s'inscrivent à temps en troisième année. Ainsi, les effets nuls sur la persévérance et l'obtention du diplôme ne peuvent pas être expliqués par la prise de conscience tardive des étudiant.e.s de leur éligibilité.

De plus, je ne trouve aucune preuve que l'éligibilité à l'aide financière supplémentaire ait eu un effet sur le type ou la qualité (mesurée par la note médiane au Bac des étudiant.e.s s'inscrivant simultanément dans le programme) du diplôme poursuivi. Les effets nuls sur la qualité du diplôme restent valables pour le diplôme obtenu un an et deux ans plus tard. Ce résultat écarte l'hypothèse selon laquelle les étudiant.e.s éligibles prennent conscience trop tard de l'aide au mérite dans le processus initial d'inscription, mais une fois conscient.e.s, choisissent ensuite de changer d'orientation vers des diplômes plus sélectifs situés dans des villes plus chères.

Il n'y a pas d'impact discernable sur d'autres mesures de l'implication dans l'enseignement supérieur, telles que le nombre d'années passées dans l'enseignement supérieur ou le niveau d'études maximal atteint, ni sur des indicateurs de performance académique tels que la probabilité de s'inscrire dans un master sélectif ou la qualité du master (à nouveau mesurée par la note médiane au Bac des pairs contemporains du programme). Bien que je ne puisse pas observer directement les notes de premier cycle des étudiant.e.s, cela indique que la performance académique ne semble pas avoir été fortement influencée par l'éligibilité à l'aide au mérite. Il n'y a pas de signe clair d'effets hétérogènes selon le sexe ou le milieu socio-économique, ce qui suggère que ces résultats reflètent de véritables effets nuls et non des effets hétérogènes qui s'annulent en moyenne. Cela implique que les trajectoires des étudiant.e.s très bon.ne.s scolairement dans l'enseignement supérieur, même lorsqu'il.elle.s viennent de milieux défavorisés, ne semblent pas être particulièrement affectées par le montant de l'aide financière qu'il.elle.s reçoivent. Je trouve des preuves d'effets positifs sur la localisation géographique (Paris et les plus grandes

viles françaises), bien que l'ampleur des estimations dépende de la fenêtre d'estimation choisie.

J'exploite l'hétérogénéité entre des sous-groupes spécifiques pour étudier trois mécanismes potentiels qui pourraient sous-tendre ces effets nuls : (i) le manque d'information sur l'éligibilité à l'aide, (ii) l'effet d'éviction de l'aide financière des parents, et (iii) l'attribution de l'aide en plus des bourses sur critères sociaux.

Tout d'abord, je ne trouve aucune preuve que l'absence d'effet sur l'*inscription* pourrait être due au fait que les étudiant.e.s ne sont pas au courant de l'existence de l'aide au mérite. Comme l'aide était automatiquement accordée aux étudiant.e.s éligibles (sous réserve de leur inscription dans l'enseignement supérieur), la question de la demande de l'aide ne se pose pas. Cependant, l'aide au mérite a été introduite en même temps qu'une vaste réforme du système de bourses sur critères sociaux, et elle n'a donc peut-être pas été aussi saillante pour les étudiant.e.s que ce dernier changement. Pourtant, les effets ne sont pas plus grands pour les cohortes de Bac plus récentes, qui sont très susceptibles d'avoir été plus au courant de l'aide au mérite, ni plus grands pour les étudiant.e.s ayant plus de camarades de lycée éligibles. Cela suggère que les déficits d'information sur l'aide au mérite ne sont probablement pas une explication de l'effet nul trouvé sur l'inscription l'année du Bac, bien que, comme discuté précédemment, cela pourrait éventuellement s'expliquer par le fait que les étudiant.e.s prennent conscience de leur éligibilité tardivement dans le processus. Comme les étudiant.e.s éligibles reçoivent l'aide financière une fois inscrit.e.s, cette préoccupation ne s'applique pas pour les résultats autres que l'inscription initiale.

Deuxièmement, j'estime les effets pour les étudiant.e.s issu.e.s des familles les plus défavorisées, qui reçoivent les montants de bourses sur critères sociaux les plus élevés mais dont les familles peuvent leur donner moins que le montant de l'aide au mérite en moyenne (Grobon and Wolff, 2022). Ainsi, pour ces étudiant.e.s, même si l'aide au mérite compense entièrement le soutien financier des parents, il.elle.s seront toujours mieux loti.e.s financièrement. Il est vrai que cela ne sera pas nécessairement le cas pour les étudiant.e.s dont les parents leur donnent plus de 200 euros par mois, et pour qui l'aide au mérite pourrait théoriquement être entièrement compensée par l'effet d'éviction. Je ne trouve aucun effet pour les étudiant.e.s dont les parents ont les revenus les plus faibles, ce qui suggère que les effets nuls globaux ne sont probablement pas le résultat de l'effet d'éviction total des contributions financières des parents par le montant reçu de l'aide au mérite. Je ne peux pas exclure des interactions potentielles entre l'éligibilité et le revenu des parents qui ne passeraient pas par le canal de l'effet d'éviction, bien que l'on puisse

s'attendre à ce que s'il y avait des effets pour un sous-groupe d'étudiant.e.s, ils concerneraient très probablement les étudiant.e.s les plus défavorisé.e.s.

Enfin, je n'observe aucune preuve que les étudiant.e.s éligibles uniquement à l'exonération des frais de scolarité et aucune allocation mensuelle dans le cadre de leurs bourses sur critères sociaux présentent des réponses comportementales plus importantes à l'éligibilité à l'aide au mérite que les étudiant.e.s éligibles à des allocations mensuelles plus généreuses dans le cadre de leurs bourses sur critères sociaux. Ces résultats restent valables même lorsque l'on restreint l'analyse aux étudiant.e.s dont les revenus parentaux sont très similaires, ce qui suggère que ces différences ne sont pas simplement le résultat de revenus parentaux différents. Cela implique que les effets nuls sont probablement pas complètement dus au fait que l'aide au mérite est accordée en plus d'autres aides financières, limitant ainsi sa capacité potentielle à produire un effet.

Cette analyse des mécanismes indique que l'explication la plus probable de l'absence d'effets observée est que les étudiant.e.s défavorisé.e.s, très bon.ne.s scolairement ne sont pas des étudiant.e.s marginaux.les, au sens où leurs résultats dans l'enseignement supérieur ne dépendent pas du montant de l'aide financière à laquelle il.elle.s sont éligibles. Cela va dans le sens de plusieurs études qui constatent de manière cohérente que l'impact de l'aide financière sur les résultats de l'enseignement supérieur tend à être faible (ou nul) pour les étudiant.e.s les plus doué.e.s scolairement, tandis que les effets pour les étudiant.e.s moins doué.e.s scolairement sont importants (Goodman, 2008; Cohodes and Goodman, 2014; Fack and Grenet, 2015; Bettinger et al., 2019; Angrist et al., 2022). Ces résultats mettent en évidence d'éventuelles complémentarités entre l'aide financière et la capacité académique des étudiant.e.s. Une avenue de recherche future intéressante consisterait à étudier plus précisément comment les effets de l'aide financière varient le long de la distribution du niveau scolaire des étudiant.e.s.

References

- Acciari, Paolo, Alberto Polo, and Giovanni L. Violante**, ““And Yet, It Moves”: Intergenerational Mobility in Italy,” *American Economic Journal: Applied Economics*, July 2022, 14 (3), 118–163.
- Aguirre, Josefa and Juan Matta**, “Walking in Your Footsteps: Sibling Spillovers in Higher Education Choices,” *Economics of Education Review*, February 2021, 80, 102062.
- Altmejd, Adam**, “Inheritance of Fields of Study,” 2023.
- , **Andrés Barrios-Fernández, Marin Drlje, Joshua Goodman, Michael Hurwitz, Dejan Kovac, Christine Mulhern, Christopher Neilson, and Jonathan Smith**, “O Brother, Where Start Thou? Sibling Spillovers on College and Major Choice in Four Countries,” *The Quarterly Journal of Economics*, June 2021, 136 (3), 1831–1886.
- Altonji, Joseph G., Erica Blom, and Costas Meghir**, “Heterogeneity in Human Capital Investments: High School Curriculum, College Major, and Careers,” *Annual Review of Economics*, September 2012, 4 (1), 185–223.
- Angrist, Joshua, David Autor, and Amanda Pallais**, “Marginal Effects of Merit Aid for Low-Income Students,” *The Quarterly Journal of Economics*, April 2022, 137 (2), 1039–1090.
- Aucejo, Esteban M., Claudia Hupkau, and Jenifer Ruiz-Valenzuela**, “Where versus What: College Value-Added and Returns to Field of Study in Further Education,” *Journal of Human Resources*, October 2022, pp. 0620–10978R1.
- Barrios-Fernández, Andrés**, “Neighbors’ Effects on University Enrollment,” *American Economic Journal: Applied Economics*, July 2022, 14 (3), 30–60.
- Bechichi, Nagui, Julien Grenet, and Georgia Thebault**, “D’Admission Post-Bac à Parcoursup : Quels Effets Sur La Répartition Des Néo-Bacheliers Dans Les Formations d’enseignement Supérieur ?,” Insee Références Édition 2021, INSEE, Paris 2021.
- Bettinger, Eric, Oded Gurantz, Laura Kawano, Bruce Sacerdote, and Michael Stevens**, “The Long-Run Impacts of Financial Aid: Evidence from California’s Cal Grant,” *American Economic Journal: Economic Policy*, 2019, 11 (1), 64–94.
- Björklund, Anders and Markus Jäntti**, “Intergenerational Income Mobility in Sweden Compared to the United States,” *American Economic Review*, 1997, 87 (5), 1009–1018.
- Black, Sandra E. and Paul J. Devereux**, “Recent Developments in Intergenerational Mobility,” in “Handbook of Labor Economics,” Vol. 4, Elsevier, 2011, pp. 1487–1541.
- , **Jeffrey T. Denning, and Jesse Rothstein**, “Winners and Losers? The Effect of Gaining and Losing Access to Selective Colleges on Education and Labor Market Outcomes,” *American Economic Journal: Applied Economics*, January 2023, 15 (1), 26–67.

- Bleemer, Zachary**, “Affirmative Action, Mismatch, and Economic Mobility after California’s Proposition 209,” *The Quarterly Journal of Economics*, December 2021, 137 (1), 115–160.
- Bowles, Samuel**, “Understanding Unequal Economic Opportunity,” *The American Economic Review*, 1973, 63 (2), 346–356.
- Campbell, Stuart, Lindsey Macmillan, Richard Murphy, and Gill Wyness**, “Matching in the Dark? Inequalities in Student to Degree Match,” *Journal of Labor Economics*, October 2022, 40 (4), 807–850.
- , –, –, and –, “Matching in the Dark? Inequalities in Student to Degree Match,” *Journal of Labor Economics*, 2023.
- Carrell, Scott and Bruce Sacerdote**, “Why Do College-Going Interventions Work?,” *American Economic Journal: Applied Economics*, July 2017, 9 (3), 124–151.
- Chetty, Raj, David J Deming, and John N Friedman**, “Diversifying Society’s Leaders? The Determinants and Causal Effects of Admission to Highly Selective Private Colleges,” NBER Working Paper 31492, National Bureau of Economic Research 2023.
- , **Nathaniel Hendren, Patrick Kline, and Emmanuel Saez**, “Where Is the Land of Opportunity? The Geography of Intergenerational Mobility in the United States,” NBER Working Paper 19843, National Bureau of Economic Research, Cambridge, MA 2014.
- , –, –, and –, “Where Is the Land of Opportunity? The Geography of Intergenerational Mobility in the United States,” *The Quarterly Journal of Economics*, 2014, 129 (4), 1553–1623.
- Cohodes, Sarah R. and Joshua S. Goodman**, “Merit Aid, College Quality, and College Completion: Massachusetts’ Adams Scholarship as an In-Kind Subsidy,” *American Economic Journal: Applied Economics*, October 2014, 6 (4), 251–285.
- Corak, Miles**, “Income Inequality, Equality of Opportunity, and Intergenerational Mobility,” *Journal of Economic Perspectives*, 2013, 27 (3), 79–102.
- , “The Canadian Geography of Intergenerational Income Mobility,” *The Economic Journal*, 2020, 130 (631), 2134–2174.
- Cortes-Orihuela, Javier, Juan Díaz, Pablo Gutierrez Cubillos, Pablo Troncoso, and Gabriel Villarroel**, “How Much Should We Trust TSTSLS Intergenerational Mobility Estimates?: Evidence From A Developing Country,” *SSRN Electronic Journal*, 2022.
- Crawford, Claire, Paul Gregg, Lindsey Macmillan, Anna Vignoles, and Gill Wyness**, “Higher Education, Career Opportunities, and Intergenerational Inequality,” *Oxford Review of Economic Policy*, 2016, 32 (4), 553–575.
- Deutscher, Nathan and Bhashkar Mazumder**, “Intergenerational Mobility across Australia and the Stability of Regional Estimates,” *Labour Economics*, 2020, 66, 101861.

- **and** –, “Measuring Intergenerational Income Mobility: A Synthesis of Approaches,” *Journal of Economic Literature*, September 2023, 61 (3), 988–1036.
- Dynarski, Susan, Cj Libassi, Katherine Micheltore, and Stephanie Owen**, “Closing the Gap: The Effect of Reducing Complexity and Uncertainty in College Pricing on the Choices of Low-Income Students,” *American Economic Review*, June 2021, 111 (6), 1721–1756.
- Estrada, Ricardo, Jeremie Gignoux, and Agustina Hatrick**, “Learning About Opportunity: Spillovers of Elite School Admissions in Peru,” 2022.
- Fack, Gabrielle and Julien Grenet**, “Improving College Access and Success for Low-Income Students: Evidence from a Large Need-Based Grant Program,” *American Economic Journal: Applied Economics*, April 2015, 7 (2), 1–34.
- Goodman, Joshua**, “Who Merits Financial Aid?: Massachusetts’ Adams Scholarship,” *Journal of Public Economics*, October 2008, 92 (10-11), 2121–2131.
- Grobon, Sébastien and François-Charles Wolff**, “Do Public Scholarships Crowd Out Parental Transfers ? Evidence from France,” *SSRN Electronic Journal*, 2022.
- Haider, Steven and Gary Solon**, “Life-Cycle Variation in the Association between Current and Lifetime Earnings,” *American Economic Review*, 2006, 96 (4), 1308–1320.
- Hakimov, Rustamdjan, Renke Schmacker, and Camille Terrier**, “Confidence and College Applications: Evidence from a Randomized Intervention,” 2022.
- Hastings, Justine, Christopher Neilson, and Seth Zimmerman**, “Are Some Degrees Worth More than Others? Evidence from College Admission Cutoffs in Chile,” Technical Report w19241, National Bureau of Economic Research, Cambridge, MA July 2013.
- Heidrich, Stefanie**, “Intergenerational Mobility in Sweden: A Regional Perspective,” *Journal of Population Economics*, 2017, 30 (4), 1241–1280.
- Hoxby, Caroline and Christopher Avery**, “The Missing “One-Offs”: The Hidden Supply of High-Achieving, Low-Income Students,” *Brookings Papers on Economic Activity*, 2013, (Spring).
- **and Sarah Turner**, “Expanding College Opportunities for High-Achieving, Low Income Students,” SIEPR Discussion Paper 12-014, Stanford Institute for Economic Policy Research, Stanford, CA 2013.
- Hoxby, Caroline M and Sarah Turner**, “Informing Students about Their College Options: A Proposal for Broadening the Expanding College Opportunities Project,” Technical Report 2013-03, The Hamilton Project 2013.
- Jacome, Elisa, Ilyana Kuziemko, and Suresh Naidu**, “Mobility for All: Representative Intergenerational Mobility Estimates over the 20th Century,” 2023.

Jerrim, John, Álvaro Choi, and Rosa Simancas, “Two-Sample Two-Stage Least Squares (TSTLS) Estimates of Earnings Mobility: How Consistent Are They?,” *Survey Research Methods*, 2016, 10 (2), 85–102.

Kirkeboen, Lars J., Edwin Leuven, and Magne Mogstad, “Field of Study, Earnings, and Self-Selection,” *The Quarterly Journal of Economics*, August 2016, 131 (3), 1057–1111.

Lefranc, Arnaud, “Intergenerational Earnings Persistence and Economic Inequality in the Long Run: Evidence from French Cohorts, 1931-75,” *Economica*, 2018, 85 (340), 808–845.

– **and Alain Trannoy**, “Intergenerational Earnings Mobility in France: Is France More Mobile than the US?,” *Annales d'Économie et de Statistique*, 2005, (78), 57–77.

Nybom, Martin and Jan Stuhler, “Biases in Standard Measures of Intergenerational Income Dependence,” *Journal of Human Resources*, 2017, 52 (3), 800–825.